UNIVERSITY OF OKLAHOMA
GRADUATE COLLEGE

GEOMETRY AGNOSTIC MUTUAL COUPLING CALIBRATION FOR PHASED ARRAY RADAR SYSTEMS

A THESIS
SUBMITTED TO THE GRADUATE FACULTY
in partial fulfillment of the requirements for the
Degree of
MASTER OF SCIENCE

By
ZACHARY SASSER
Norman, Oklahoma
2025

GEOMETRY AGNOSTIC MUTUAL COUPLING CALIBRATION FOR PHASED ARRAY RADAR SYSTEMS

A THESIS APPROVED FOR THE
SCHOOL OF ELECTRICAL AND COMPUTER ENGINEERING

BY THE COMMITTEE CONSISTING OF

Dr. Caleb Fulton, Chair

Dr. Mark Yeary

Dr. Clifford Fitzmorris

Table of Contents

List of Tables

List of Figures

Abstract

Phased array radar systems require precise calibration to ensure accurate beam steering and signal processing, as any variations in the individual antenna elements can significantly compromise performance. Traditional calibration methods, often reliant on known geometries and specialized equipment, face limitations in real-world applications where data sets may be incomplete due to engineering constraints, such as amplifier saturation and varying element configurations. This thesis evaluates existing calibration algorithms, highlighting their dependency on geometric patterns and the complexities involved in handling missing data.

To address these challenges, we propose a geometry-agnostic calibration methodology that utilizes mutual coupling measurements to achieve robust calibration without reliance on specific geometrical configurations. By leveraging various computational techniques, our approach offers the flexibility needed for in-situ and initial calibration. This methodology not only generalizes previous calibration techniques but also enables applications in radar systems utilizing non-planar and independently deployed antennas. The findings suggest that the proposed method can significantly enhance calibration accuracy and operational efficiency, opening new avenues for radar applications where traditional methods are impractical.

Chapter 1

Introduction

Phased array radar systems require calibration to ensure accurate beam steering and signal processing. These systems rely on the precise control of the phase and amplitude of signals transmitted or received by an array of antennas. Calibration corrects for variations in the individual antenna elements and the associated electronics, which can cause errors in the direction and shape of the radar beam. Without proper calibration, the radar's performance in detecting, tracking, and imaging targets would be compromised. A common way of achieving calibration, particularly for stationary systems, is to use an external calibration antenna to align each element. This generally requires constructing a near-field scanner that is large enough to contain the radar panel. However, in more recent years, algorithms have been developed to achieve calibration through mutual couplings [1]–[12]. The data collection process involves transmitting from one element and then measuring the resulting signal for each receiving radiator. The peak of the matched filter's response of the transmitted signal is then divided by the peak of the matched filter's response of the received signal to obtain a complex coupling coefficient. The data are then collected into a coupling matrix which can be used alongside other reference information to compute the alignment weights for each element.

Phased array calibration algorithms can be classified by various metrics. One major distinction is whether the algorithm is to be applied in the laboratory or factory setting in which an operator will generally have access to specialized equipment such as near-field scanners, far-field scanners, or a simple reference antenna with a known complex gain. These calibration algorithms are called "initial calibration" algorithms and while they are generally more precise, they are often impractical for some radar systems and cannot be used for calibration after the radar is deployed. Calibration algorithms that are designed to be used while the radar is deployed in the field are called "in-situ" calibration algorithms. Another defining attribute of a calibration algorithm is its dependence on external equipment or built-in equipment such as a calibration antenna or a coupled line at each element. Recent developments have been made towards algorithms that do not rely on additional hardware in order to achieve calibration using mutual coupling measurements [1]–[8], [10]–[12].

In this research, we shall evaluate many of the existing mutual coupling calibration methods and discuss their particular drawbacks. The common drawback of the algorithms to be discussed is their dependence on known geometry which often requires hard-coded logic to solve effectively [4], [6], [10]. This poses challenges when implementing these algorithms on real-world systems as the mutual coupling data set is often incomplete due to various engineering constraints that arise when designing an RF front-end. For example, elements that are geometrically close together will often cause amplifier saturation which introduces non-linear behavior to the system [4], [13]–[15]. It may at first appear that the solution is to reduce the transmit power until the receiving element is operating below saturation. However, this will limit the range of the radiated power, resulting in further away elements having

a receive power that is too low to be a meaningful measurement. An ideal algorithm will have the flexibility to automatically solve around missing data, whether that be from a faulty channel, missing panel, amplifier saturation, etc. Algorithms that are heavily dependent on a particular geometry often require a significant amount of extra logic to handle missing data.

A geometry-agnostic methodology is proposed in this thesis alongside the analysis of the existing mutual coupling calibration methods. In addition to having the capability of theoretically calibrating any arbitrary configuration of antennas, the process is intended to support the previous methods of mutual coupling calibration. It is now feasible to simulate estimated passive couplings through methods like method of moments [13], [15] or using software like Ansys Electronics Desktop (HFSS) [16]. Instead of making geometric assumptions about the passive couplings between elements, we shall evaluate the couplings directly based on our estimate. If used as an in-situ calibration, the passive couplings can be obtained by taking a mutual coupling measurement after the array has been calibrated by an initial calibration method. In the in-situ case, the passive coupling matrix is not estimated but directly measured. If the guided approaches proposed here are used, it is not even necessary for the passive coupling estimate to be numerically accurate. In those cases, the passive couplings are only used as a guide to determine which pairings are close. If a coupling is approximately equal to another coupling when the array is calibrated, they should also be close in the passive coupling estimate; the passive couplings do not necessarily need to accurately predict the calibrated couplings. An advantage of this methodology is the generalization that enables previous methods to be used in the same underlying computational framework. The next section discusses the passive coupling matrix in more detail.

1.1 Passive Couplings ("reference matrix")

Two types of coupling matrices for mutual coupling calibration exist. First, there is the passive coupling matrix as shown in Fig. 1.1. In prior literature, this was simply referred to as S-parameters, this terminology would not be accurate in all cases. For example, couplings include the effects of load and source impedance while S-parameters do not. The way in which the matrix is used will be mostly the same. The passive coupling is a complex number that represents how the signal is changed from transmission to receive when the gains from both of the frontends are normalized to unity gain. Generally, this is intended to characterize the effects of electromagnetic transmission between the two radiators imposed on the signal. It is important to note that passive coupling is a function of both frequency and environment.

1.2 Active Couplings ("measured matrix")

Active couplings are a function of the passive couplings and are a characterization of the total change to the signal, shown in Fig 1.2. The complex gains characterize the amplifiers and any other effects imposed on the signal by the front-end [11], [12]. A visual example of how this data is represented is shown in Fig. 1.3. A mutual coupling measurement should be at the same frequency and in a comparable environment to that which was used for the passive coupling. The ultimate goal of calibration is to compute gains so alignment weights can be applied to the amplifiers aligning each radiator in magnitude and phase. If the system is linearized, the alignment weights should map the passive coupling matrix to the active coupling matrix.

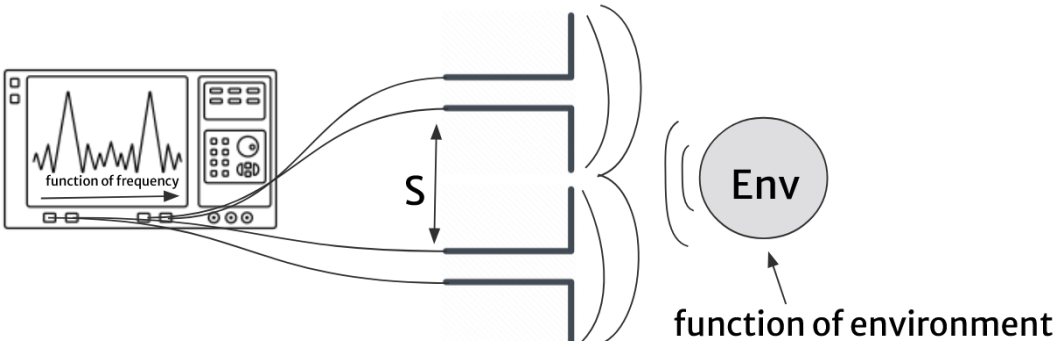

Figure 1.1: An illustration to convey the dependence of passive couplings on frequency and environment. A VNA is shown to represent the passive couplings because in many circumstances the passive couplings are equivalent to the S-parameters. A metal sphere is shown to illustrate that the environment the passive coupling is measured in is relevant.

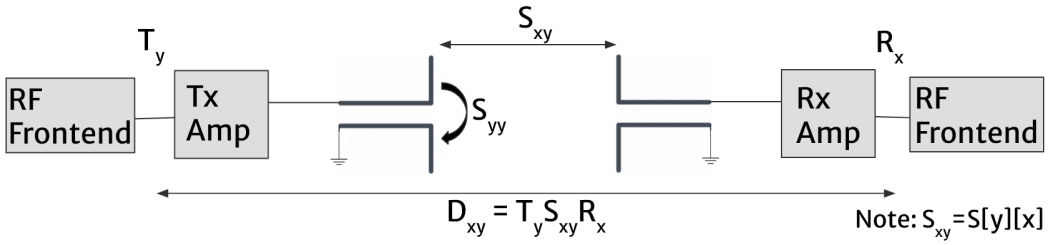

Figure 1.2: A signal path block diagram illustrating that active couplings (measurement matrix) are directly related to the passive couplings but includes the effects introduced by the amplifier, line geometries, and other associated frontend contributions by representing the overall effect with a complex value known as a gain. There is a gain associated with both transmit and receive.

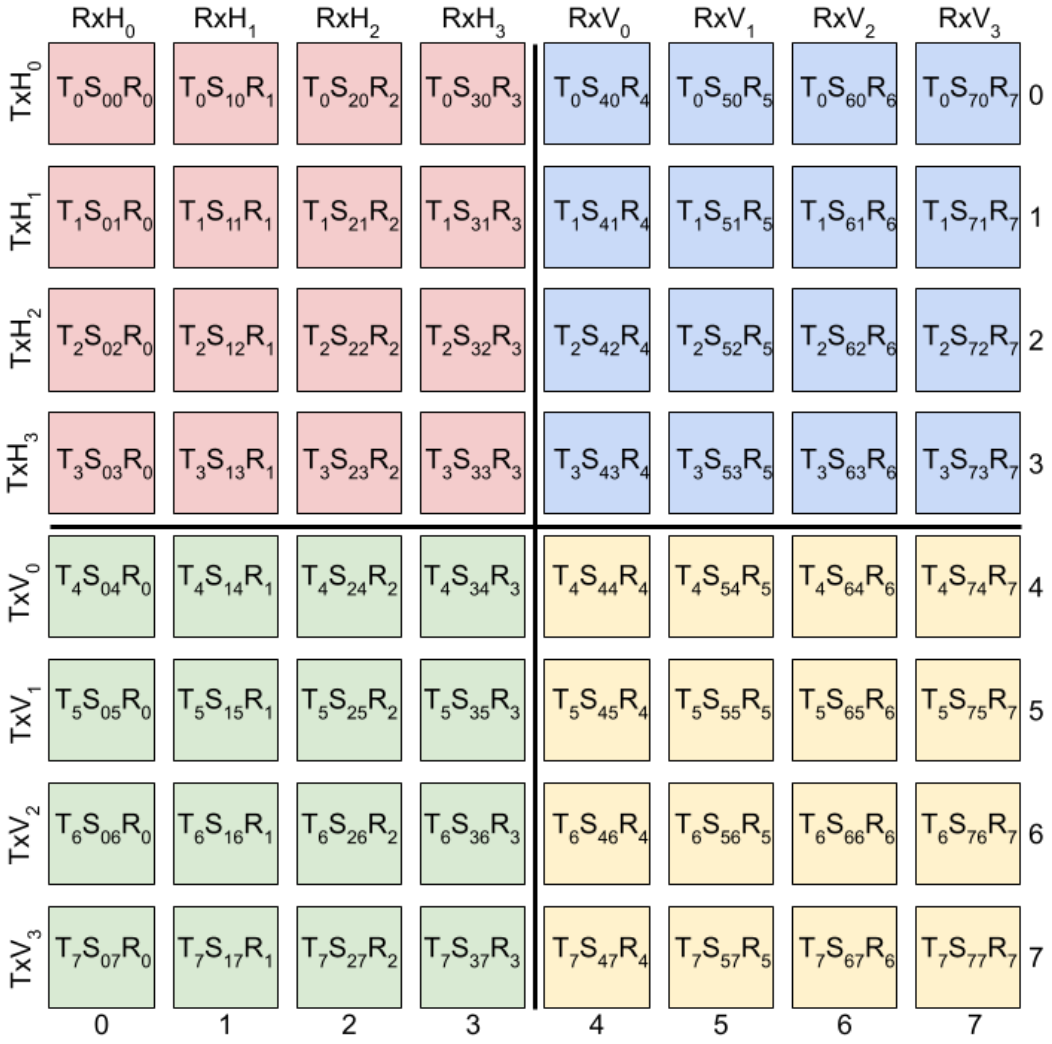

Figure 1.3: A visual depiction of the structure of the coupling matrix used to implement these algorithms. The example shown is a 4-element dual pol system. The coupling matrix can be seen as a multiplication of a column vector containing the complex transmit gains, multiplied by a row vector containing the complex receive weights, and then element-wise multiplying the passive coupling matrix. It is important to remember that to index different polarizations (radiators on the same element), a multiple of the number of radiators must be added to its element index. This is different from the format output by some radar systems, such as the 6D representation ($P_{Tx},\ X_{Tx},\ Y_{Tx},\ P_{Rx},\ X_{Rx},\ Y_{Rx}$).

1.3 Estimation of Passive Couplings

With modern computing power, it is now feasible to simulate passive couplings in a variety of ways. The simulation method to use is dependent on the application and accuracy required. As input for the methods we will describe, it is assumed that an estimate of the passive couplings is provided. In some cases it may not be necessary for this data to be perfectly accurate, of course, the better the estimate, the better the results. If it is feasible to do an initial calibration using a higher accuracy method, such as using advanced equipment like far field and near field scanners, then the mutual coupling methods described here are better suited as an in-situ algorithm. Instead of estimating the passive couplings, a coupling measurement is taken while the system is in a calibrated state. This coupling measurement is referred to as the "reference matrix". A perfect estimation of passive couplings would be equal to this reference matrix. The reference matrix must be taken in a similar environment and at the same frequency as the couplings that will be taken in the field. Generally, it is best for both measurements to be taken in an environment similar to free space. If multiple frequency ranges are to be used, it is necessary to obtain different weights for different frequencies. Broad frequency ranges require a more complex system than simple weights. As the frequency moves from the frequency in which the weights were generated, the error increases. This could potentially be handled by swapping sets of alignment weights rapidly as frequency changes. Another potential option is to apply weights as functions of frequency rather than as constants. The function could be composed by using linear interpolation between alignment weights obtained at various frequencies. There is much potential for further research on these topics.

1.4 Known Reference Gain

For most mutual coupling calibration algorithms, it is necessary to provide a known reference gain for at least one radiator (ideally for both transmit and receive). Note: for a dual-pole system, each element has 2 radiators. For the geometry agnostic methods that will be proposed in this thesis, the algorithm treats every radiator as unique. If this reference information is not provided, the system of equations will always be circularly defined. Obtaining relative gains without a reference is possible, which can be useful depending on how alignment weights are applied in the frontend systems. In the cases where a reference gain cannot be provided, every radiator can be aligned to a reference radiator, which for many applications may be adequate. The error will manifest as a signal delay, but one that is constant across all elements since they are aligned to each other. In this situation, there is also an extra magnitude scaling step that must be performed by comparing the magnitudes of the reference matrix to the measurement matrix. Whether it is necessary to provide a reference gain is dependent on the application, however, generally it is best to provide some kind of estimate to constrain the solution space. It is also very simple to include multiple reference gains to improve accuracy as will be shown later.

1.5 Solution Range

An infinite number of potential alignment states can be achieved with any alignment method. This is because any non-unity change imposed by the transmit gain can be countered by an equal and opposite change imposed by the receive gain as can be seen in eq. (1.1). This leaves a lot of room for magnitude and phase ambiguity. For example, if all transmitters are aligned to a phase of 30 degrees with a magnitude of one, adequate alignment can still be achieved with receivers that are aligned to a phase of -30 degrees. With phase alone, there exists an infinite number of potential alignment states. To obtain another solution one can simply apply a rotation to the transmit weights along with the opposite rotation to the receive weights. The same is partially true for magnitude effects in theory, but in practice, magnitude cannot have the same level of variation. For example, if the transmit gains are aligned to a magnitude of 3 dB, the receive gains can be aligned to counter the amplified signal by aligning to a magnitude of -3 dB. However, real-world hardware considerations will severely limit the extent to which magnitudes can change because effects such as amplifier saturation will begin to take place. Similarly, the received signal is always orders of magnitude weaker than the transmit power. Magnitude offsets introduce the potential for receiving signals which are too small to obtain accurate data.

$$T \cdot R = 1 \implies T = R^{-1} \tag{1.1}$$

where $T, R \in \mathbb{C}$

1.6 Drawbacks of Existing Methods

Datasets collected on real-world radar systems are often not complete enough to implement the currently existing methods of mutual coupling calibration. For example, there is a range limitation for elements that are far apart. For low transmit power, the pairings that are geometrically far away become less accurate or don't exist at all due to measurement limitations. However, it is not feasible to simply increase the transmit power as elements that are close together will cause the receiver LNA to saturate. While receiver saturation is more prevalent for a high transmit power measurement, these effects are often still prevalent in the low power case. While these effects can be minimized by careful consideration of transmit and receive amplifier settings, at the time of writing it is simply not feasible to eliminate both of these effects. There are many other factors that may also result in missing data, and for a calibration algorithm to be robust it must be capable of handling arbitrary missing data.

The critical drawback of these methods of calibration is their reliance on geometric patterns. The existing algorithms are often reliant on hard-coded logic, which means that when the program encounters a piece of missing data, there must be extra logic to reroute the process so that the information that is missing can be obtained elsewhere. For example, the lattice method described in the next chapter becomes exceedingly difficult if even one pairing in the chain cannot be solved. One solution is to increase the size of the lattice when these conditions occur, however, this may require complex logic as it may be necessary to create a larger gap lattice to solve a few elements and then scale back down to a smaller gap lattice to solve the rest. This is even more problematic considering the fact that the 1-element spacings will often

not be available in the dataset due to the saturation problem. While larger lattices can be used, the logic becomes increasingly complex and in some cases unsolvable.

The coupling set methods described in [4], [6], [10] are also dependent on repeating patterns that require similarly complex logic to avoid missing data. In addition to this, the base assumption made by the coupling set method is very often not accurate. Due to a variety of factors such as edge effects, similarly spaced elements do not necessarily have the same passive coupling. The statistical rejection of certain sets does help to eliminate these effects, but it is not a perfect solution [12], [17]. There may exist much more adequate pairings (closer passive couplings and more linearly independent equations) within the dataset that may not follow an obvious geometric pattern. By only considering pairings for which we can define a geometric pattern, we may be missing valuable combinations of pairings that would improve the solution accuracy.

Finally, the reliance on geometric patterns restricts the radiator configuration to geometries that are simple to create indexing patterns. Non-planar and unequally spaced radiators have a large potential for radar applications such as passive radar or applications where antennas may be deployed independently around the environment. These applications would make existing calibration methods very difficult or even impossible. A traditional park and probe initial calibration ("factory calibration") might be impossible and/or impractical for a radar system that operates on multiple field deployable antennas that do not have fixed locations or that are too large for a scanner. There are a variety of future and current applications that could benefit from geometry-agnostic calibration.

Chapter 2

Pre-alignment

2.1 Phase Wrapping

Before we can feed a calibration algorithm with data, we must deal with the condition of phase wrapping. As described in [4], [10], the application of logarithms discussed in a later section imposes the constraint that the phase variation of the Tx and Rx gains must be within $\pm 90^\circ$ (see Appendix B). To address this, we first apply a phase alignment algorithm to determine alignment coefficients that correct for phase wrapping effects. The overall adjustment weights applied to the system will be the phase alignment weights multiplied by the weights determined by the alignment process itself. It is important to understand that this is simply a pre-phasing step and not the actual alignment algorithm. Although, in some cases, alignment in phase can be achieved with only pre-phasing. The ideal case would result in perfect alignment of the elements in phase; however, this is not strictly necessary for the alignment algorithm to produce quality calibration results. In the next section, we will discuss is an example which was first described by [18] that will highlight geometric dependency and vulnerability to missing data.

2.2 Lattice Phase Alignment

In this section we will apply a method which falls under patent [18] for demonstration. First, we must form lattice structures, as shown in Fig. 2.1. To align every element, four lattices must be used (more if larger spacings are used). The phase alignment algorithm begins by aligning the four reference elements in the center of the array to each other. This ensures that the lattices extending from each reference are aligned to each other. The process involves two key steps: (1) compensating for the phase difference caused by unequal distances between elements, and (2) calculating the remaining phase offset using active coupling measurements.

Initial Alignment of Reference Elements

The initial alignment of the four reference elements is critical to ensure that the lattices are properly aligned. Consider two reference elements, ref(1) and ref(2), and a transmitter T. The phase difference between the two reference elements due to their unequal distances from T can be expressed as:

$$\theta = \frac{2\pi f \cdot d}{c},$$

where:

- f is the frequency of operation,
- d is the distance difference between T and the two reference elements,
- c is the speed of light.

This phase difference must be removed to isolate the true phase offset between the reference elements. Let $D_{T,\text{ref(1)}}$ and $D_{T,\text{ref(2)}}$ be the active coupling measurements between T and the two reference elements. The ratio of these couplings is given by:

$$K = \frac{D_{T,\text{ref(1)}}}{D_{T,\text{ref(2)}}}.$$

The phase difference caused by the unequal distances is removed by dividing K by $e^{j\theta}$. The remaining phase difference, $\Delta\phi$, is the true phase offset between the two reference elements:

$$\Delta\phi = \angle\frac{K}{e^{j\theta}}.$$

This phase offset is then applied to align ref(2) to ref(1):

$$\text{RxOffset}(\text{ref}(2)) = e^{j\Delta\phi}.$$

The corresponding row of the coupling matrix D is adjusted accordingly:

$$D(:, \text{ref}(2)) = D(:, \text{ref}(2)) \cdot \text{RxOffset}(\text{ref}(2)).$$

This process is repeated for all pairs of reference elements to ensure that they are aligned to each other. Once the reference elements are aligned, the lattices extending from them can be aligned using the same method.

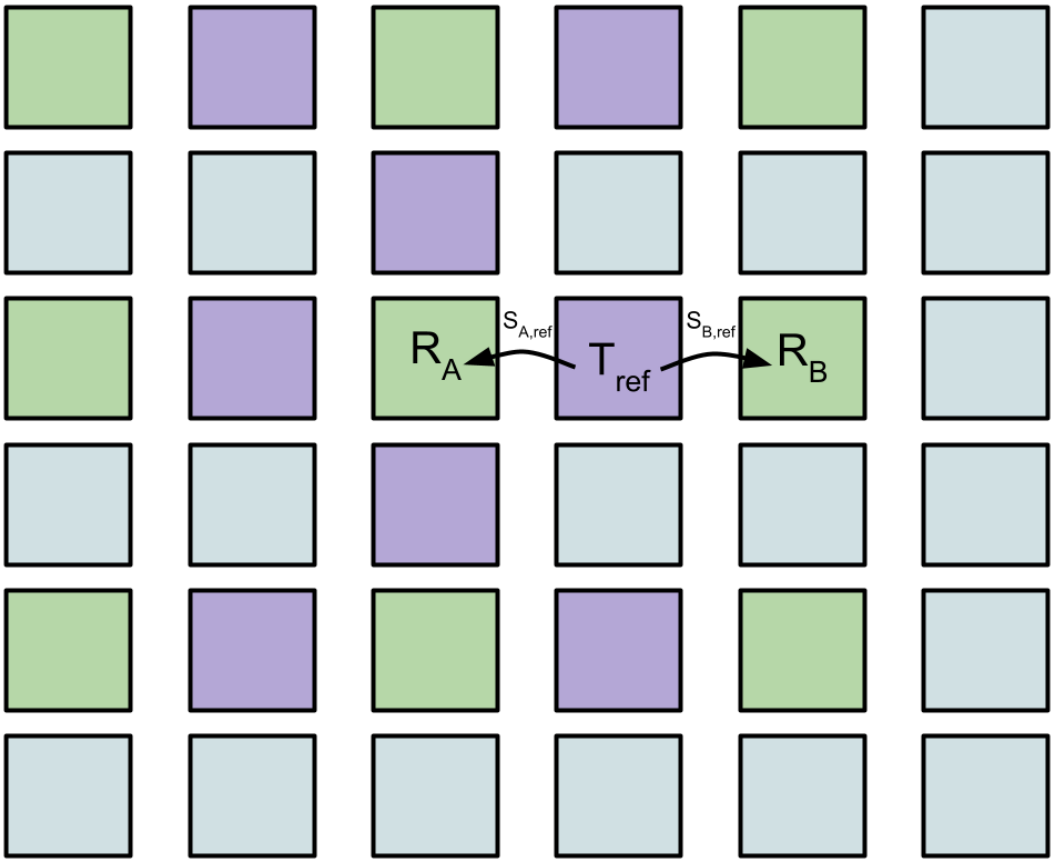

Figure 2.1: One of the four lattice structures used for the lattice phase alignment method. The elements highlighted in purple are used as references to align the elements shown in green to each side of the particular reference element. Only one pairing set is labeled for simplicity. After the process is complete, every element marked green will be aligned to the element in the center of its lattice. The process is applied in exactly the same manner to align transmit. The remaining elements will be aligned by the other three lattices.

Alignment of Lattice Elements

After aligning the reference elements, we proceed to align the remaining elements in the lattice. Using the aligned reference elements as a starting point, we calculate the phase offsets for the surrounding elements. This is done by applying (2.1) from [6]–[8], [12], which uses two active coupling measurements with a common reference transmitter (T_{ref}) coupling with receivers on opposite sides (R_A and R_B), as shown in Fig. 2.1. Due to the symmetry, the passive coupling values ($S_{A,\text{ref}}$ and $S_{B,\text{ref}}$) can be assumed to be approximately equal. This implies that the division of the two uncalibrated active couplings ($D_{A,\text{ref}}$ and $D_{B,\text{ref}}$) results in an expression that describes the ratio of the receive gain of radiators R_A and R_B:

$$\begin{aligned} K_{ref,B}^{ref,A} &= \frac{D_{A,ref}}{D_{B,ref}} = \frac{T_{ref}S_{A,ref}R_A}{T_{ref}S_{B,ref}R_B}, \\ S_{A,ref} \approx S_{B,ref} &\Rightarrow K_{B,ref}^{A,ref} = \frac{T_{ref}R_A}{T_{ref}R_B} = \frac{R_A}{R_B}, \\ \therefore \angle K_{ref,B}^{ref,A} &= \angle R_A - \angle R_B = \triangle\phi. \end{aligned} \tag{2.1}$$

By adjusting the phase of one of the radiators by $\Delta\phi$, the receive gains are aligned in phase. This process is repeated for all pairs of elements in the lattice, working outward from the reference elements. The coupling matrix is also adjusted with the phase alignment weights before proceeding to the alignment algorithm. One possible resulting alignment state is shown in Fig. 2.2. The variation in the results is caused by the approximation used to align the lattices to each other, as well as a small amount of measurement error from real-world conditions.

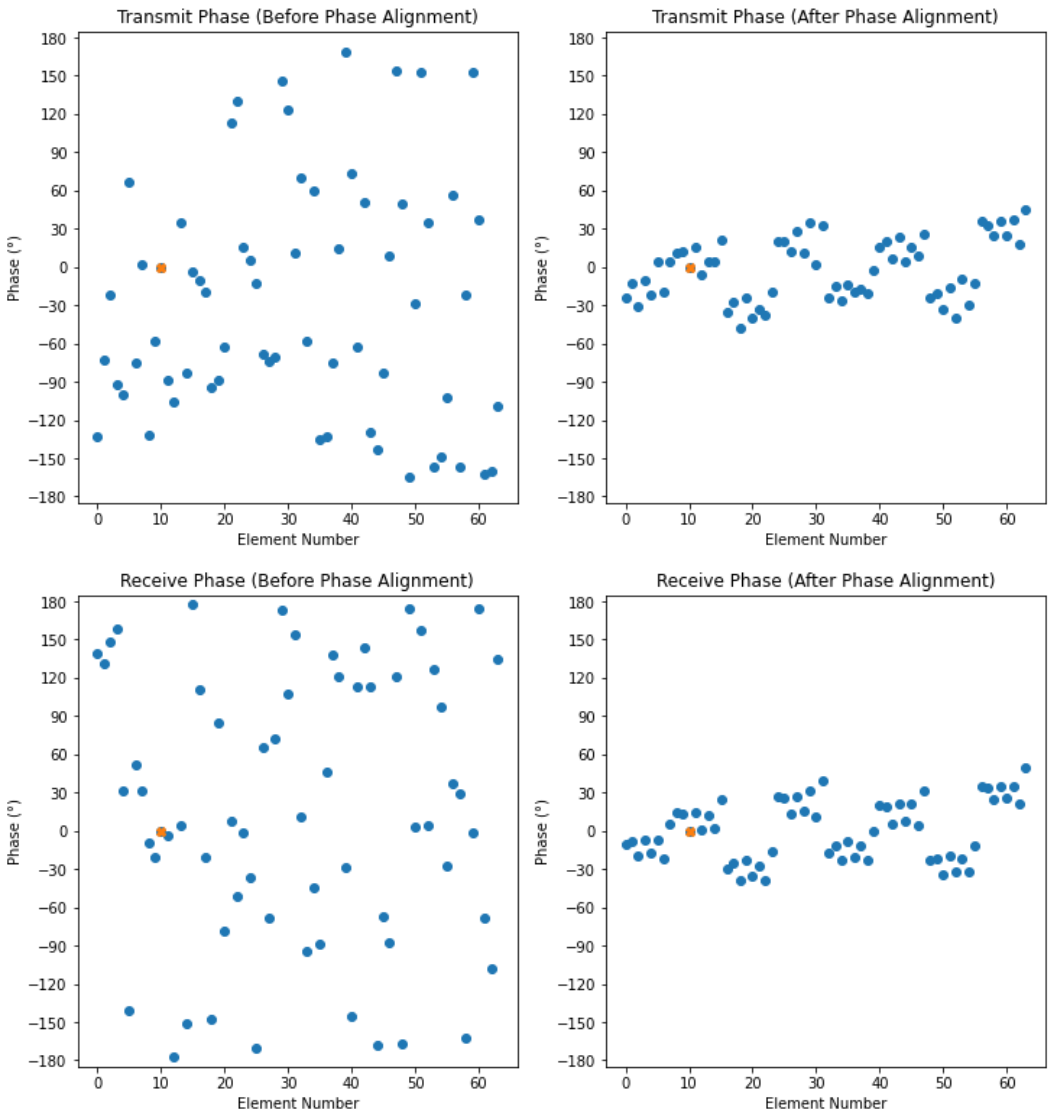

Figure 2.2: The result of lattice alignment applied to a dataset gathered on an 8x8 phased array system. After lattice alignment has been completed, the phases of each gain are within an acceptable range to avoid phase ambiguity. The four colored dots represent the center element of each lattice. The variation is introduced by the approximation step used to align the lattice centers, as well as a small amount of measurement error from real-world conditions.

2.3 Simple Alignment

If the passive couplings are known with perfect accuracy, it is possible to do a trivial alignment. The process is shown in eq. (2.2) through (2.9) after the coupling data has been divided by the appropriate passive coupling. The variables *a,b,c,d,e,* and f are used to express the value of the gains before each equation is applied (since T_1 through R_3 will change as the weights are applied in each step). In eq. (2.2), we arbitrarily choose element 1 to act as the reference. We then pick a pairing involving transmitting from element 1 to element 2 to compute an alignment weight for the receiver of element 2. The coupling data already has the passive coupling effects removed, so the coupling represents the multiplication of gains, ae. If we then adjust the gain of R_2 by $1/D_{21} = 1/ae$, the gain of R_2 will then be $1/a$ as the e term will cancel. Now that R_2 has been aligned to a relatively known value, we can then use it to align another transmitter.

$$Let,\ \{T_1 = a,\ T_2 = b,\ T_3 = c,\ R_1 = d,\ R_2 = e,\ R_3 = f\} \in \mathbb{C}$$

$$T_1 R_2 = D_{21} = ae \tag{2.2}$$

$$Adjust\ R_2\ by\ \frac{1}{D_{21}} = \frac{1}{ae}$$

Using the same logic as before but for a pairing from element 3 to 2, T_3 can be adjusted by $D_{21}/D_{23} = a/c$ which will change the gain of T_3 from c to a. Note that the alignment weight for R_2 is included in the expression, which

is why the adjustment includes a D_{21} term. If the coupling matrix is updated with each step, it is unnecessary to include the prior weights in the calculation as D_{23} will already be c/a instead of ce. The process then continues until every element is solved. The application of this algorithm is more of an exercise to help build intuition for the logic that is automated by linear algebra we will use later. Those versed in linear algebra will notice how this can be expressed more efficiently in matrix notation. It has been shown this way for deeper understanding rather than practical implementation.

$$\implies T_3R_2 = D_{23} * \frac{1}{D_{21}} = ce * \frac{1}{ae} = \frac{c}{a} \tag{2.3}$$

$$Adjust\ T_3\ by\ \frac{D_{21}}{D_{23}} = \frac{a}{c}$$

$$\implies T_3R_1 = D_{13} * (\frac{D_{21}}{D_{23}}) = cd * \frac{a}{c} = ad \tag{2.4}$$

$$Adjust\ R_1\ by\ \frac{D_{23}}{D_{13}D_{21}} = \frac{1}{ad}$$

$$\implies T_2R_1 = D_{12} * \frac{D_{23}}{D_{13}D_{21}} = bd * \frac{1}{ad} = \frac{b}{a} \tag{2.5}$$

$$Adjust\ T_2\ by\ \frac{D_{13}D_{21}}{D_{12}D_{23}} = \frac{a}{b}$$

$$\implies T_2R_3 = D_{32} * \frac{D_{13}D_{21}}{D_{12}D_{23}} = bf * \frac{a}{b} = fa \tag{2.6}$$

$$Adjust\ R_3\ by\ \frac{D_{12}D_{23}}{D_{32}D_{13}D_{21}} = \frac{1}{fa}$$

$$\implies T_1R_3 = D_{31} * \frac{D_{12}D_{23}}{D_{32}D_{13}D_{21}} = af * \frac{1}{fa} = 1 \tag{2.7}$$

$T_1\ needs\ no\ adjustment\ (it\ is\ a\ reference)$

$$\implies T_1 = a,\ T_2 = b\ *\ \frac{a}{b},\ T_3 = c\ *\ \frac{a}{c} \tag{2.8a}$$

$$R_1 = d\ *\ \frac{1}{ad},\ R_2 = e\ *\ \frac{1}{ae},\ R_3 = f\ *\ \frac{1}{fa} \tag{2.8b}$$

$$\therefore T_1 = T_2 = T_3 = a\ and\ R_1 = R_2 = R_3 = \frac{1}{a} \tag{2.9}$$

However, this method has the same drawback as the previous method, if a missing piece of data is encountered, extra logic must be implemented to find another set of pairings that can solve for the information. For the example below, if D_{21} happens to be missing or inaccurate, the solution breaks. The program reaches a state that it must decide how to get the information from other pairings and reroute the solution. The case of 3-elements is particularly sensitive as a single missing data entry makes the system unsolvable as the dataset simply has insufficient information. Note that the solution is relative to the reference element; if the reference element gain is known, the system can be calibrated to gains of unity magnitude and zero absolute phase.

After this process is applied, every transmit gain will be aligned with every other transmit gain. Likewise, every receive gain will be aligned with every receive gain. For this particular algorithm, the receive gains will always end up as the inverse of the transmit gains (if a receive reference is used, the relationship will be reversed). Note that all of the gains are aligned with reference to the reference gain T_1. If a reference gain is provided, all of the transmit gains would be divided by the known reference gain and all the receive gains would be multiplied by the known reference gain. After applying the known reference, all of the elements will be aligned to one. The primary limitation of this methodology is that the solution error is highly sensitive to the error in the passive coupling data. A robust algorithm should isolate estimated values such as the passive coupling matrix from measured data to limit the effect of our estimations on solution accuracy. By including the estimated passive coupling directly in the calculation, the error associated with that estimation directly contributes to the error of the results.

2.4 Passive Coupling Subtraction Method

The previously mentioned methods of phase alignment all rely on geometric knowledge of the system. The focus of this paper is on designing a methodology for calibration which does not require prior knowledge of the geometry of the array. Fig. 2.3 depicts the phase alignment routine that only requires an estimated or previously measured passive coupling matrix in order to determine appropriate alignment weights. We first convert our measured coupling matrix (uncalibrated) to its phase angle representation. This effectively makes every index of the matrix the sum of its corresponding transmit gain, receive gain, and passive coupling phase. From there, perform the same operation with the passive coupling matrix and then subtract the reference coupling angular matrix from the measured coupling angular matrix. At which point every index of the resulting matrix is the sum of the transmit gain phase, receive gain phase, and the error in the estimate matrix. This stage is the primary source of error for this phase alignment method as the estimated matrix is directly used in the computation. We can then choose a particular element to act as a reference by subtracting all rows of the matrix with the row corresponding to the reference element. This represents all phases as relative to the reference element. Note that it is still unknown what the particular absolute phase of the reference transmit gain was. It is only known what the phases are relative to that element. Next, we take the modulo with respect to 360 deg to avoid representing equivalent angles as different values (ex. -45 deg and 315 deg).

After this operation, the phases should be the same across a given row of the matrix. Errors in the passive coupling estimate or in the coupling measurements themselves will cause variation across the row. To limit this error and decide on a value, the mean of each row is taken. The resulting values are the offsets to adjust each transmit element. The process can then be repeated for receive by instead applying according to the columns of the matrix. Figures 2.4 through 2.7 depict an example of this process.

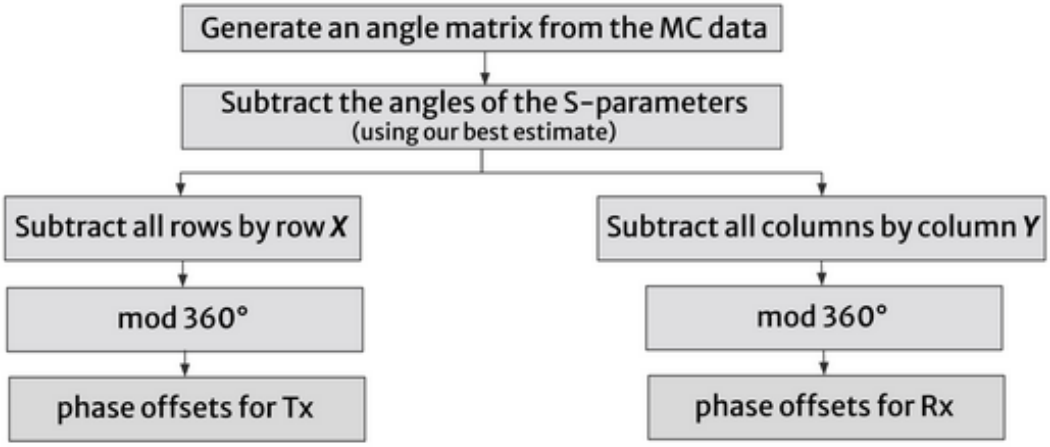

Figure 2.3: The steps involved in geometry agnostic phase alignment. Passive coupling phases are subtracted from the measured matrix phases. To obtain phase alignment weights, all rows are subtracted by the row of a reference element, resulting in all phases being in relation to the reference. A modulo with 360 degrees puts all values into a one to one representation. The mean of the row is taken to obtain the phase alignment weights for each element. The process is then repeated for receive with columns instead of rows.

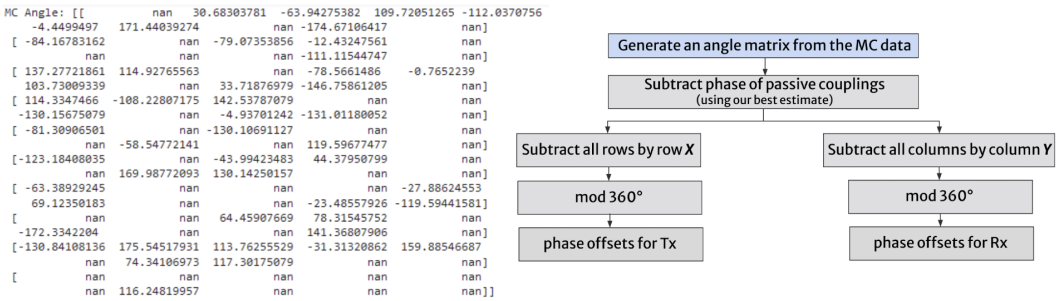

Figure 2.4: The state of the mutual coupling measurement matrix after applying the angle function to the data. At this stage the data is representing the sum of the phase of the transmit gain, receive gain, and associated passive coupling (wrapped since phase is a modular space).

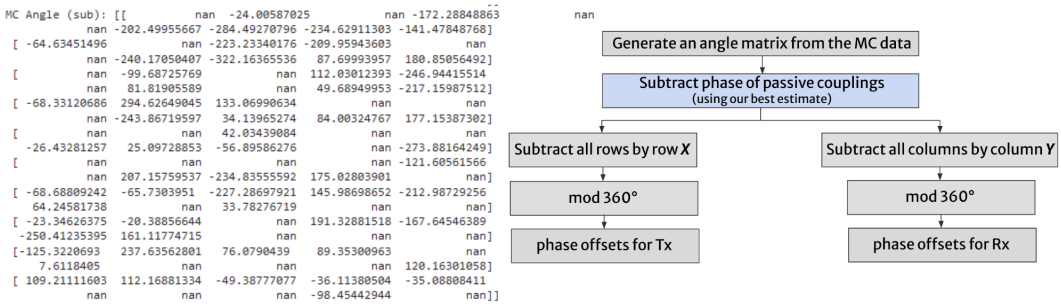

Figure 2.5: The state of the modified coupling matrix after removing the phase contribution of the associated passive coupling. At this stage, each value represents simply the sum of the phase of the receive and transmit gain, along with phase error present in the passive coupling reference matrix.

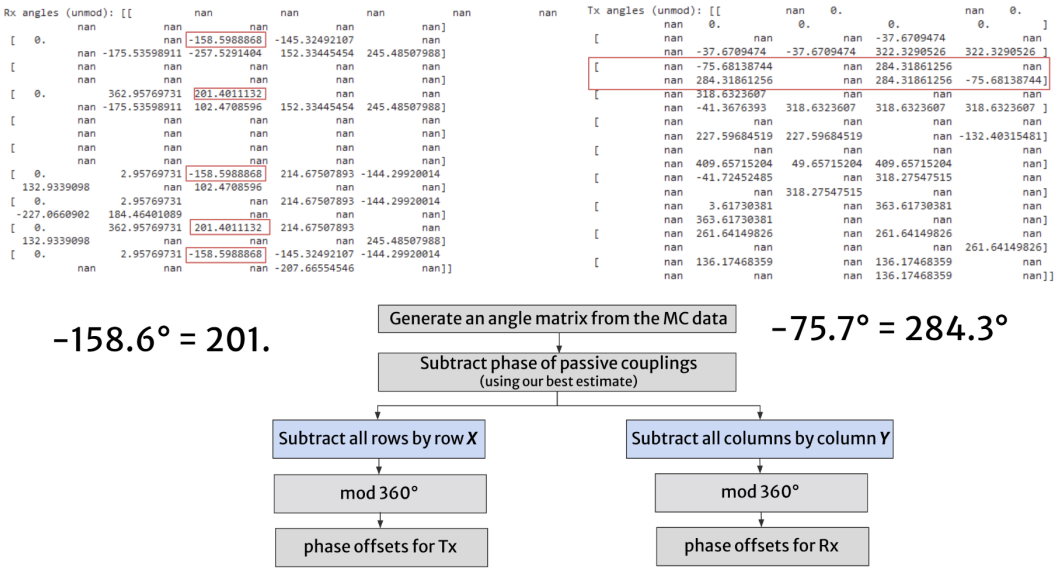

Figure 2.6: The state of the matrix after transforming the data to be in relation to the phase of a reference element. Here we can see that every row/column is actually the same phase but with some entries being the negative representation. The diagram splits here as this operation is done separately for the transmit and receive. The matrix shown on the left represents the Tx path (left) on the flowchart and the right matrix represents the Rx path (right).

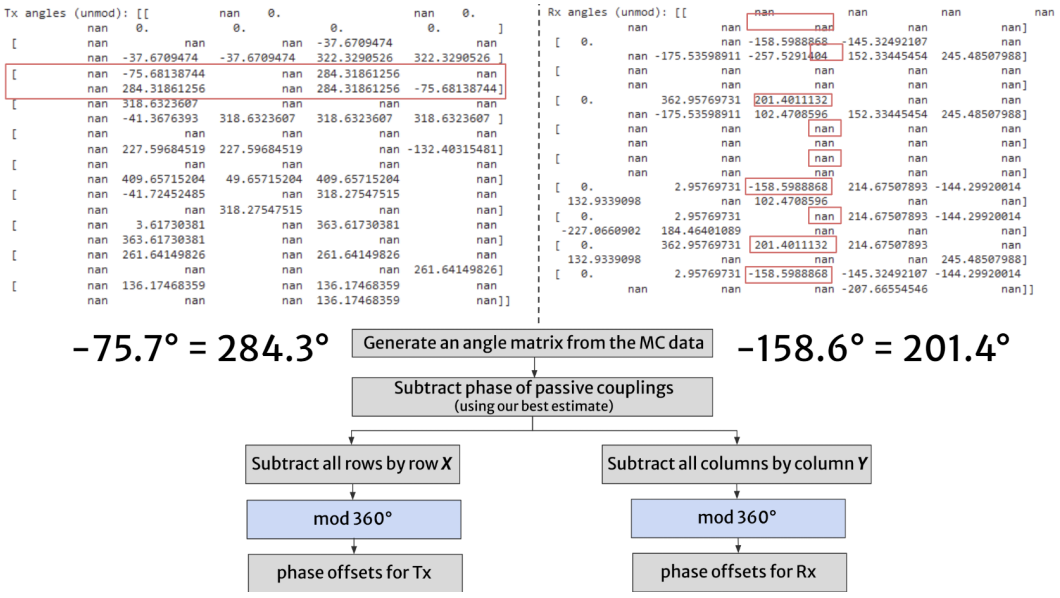

Figure 2.7: After applying a modulo of 360 deg or 2π, equal angles are represented with the same number. At this stage, we can see that the resulting rows/columns are equal. After these adjustments are applied, all gains will be aligned with their respective reference. The matrix shown on the left represents the Tx path (left) on the flowchart and the right matrix represents the Rx path (right).

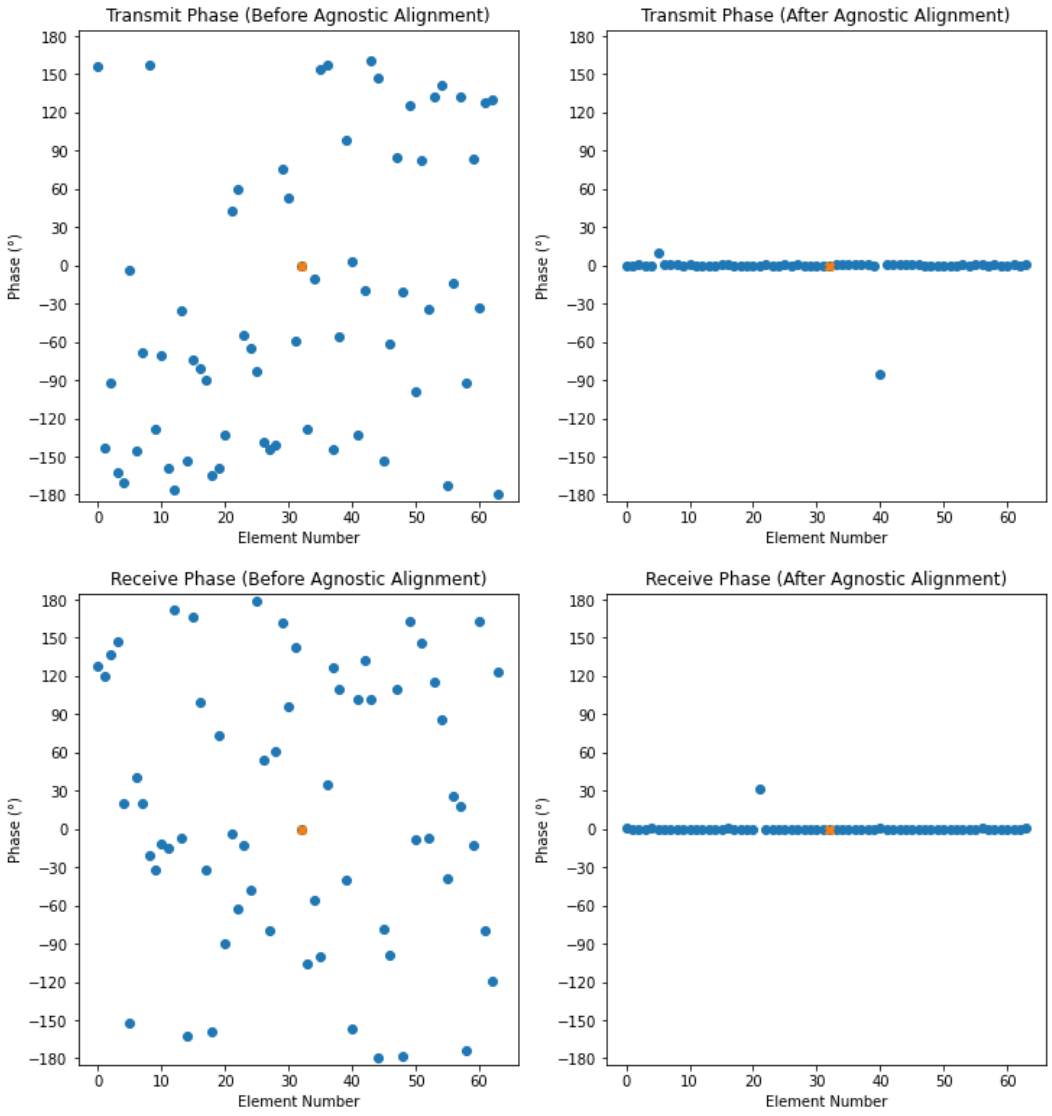

Figure 2.8: The result of subtractive alignment applied to a dataset gathered on an 8x8 system. After alignment, the phase of each gain is identical with 3 acceptable outliers. This accuracy is heavily dependent on the accuracy of the passive coupling matrix. In this case, a coupling matrix taken when the array was calibrated was used. This is feasible when applying the algorithm in-situ or if a dataset can be taken from a comparable system. If lower accuracy passive coupling datasets are used, the results will deteriorate accordingly. The matrix shown on the left represents the Tx path (left) on the flowchart and the right matrix represents the Rx path (right).

2.5 Passive Coupling Guided Phase Alignment

The subtractive method of phase alignment is directly limited by the accuracy of the provided passive coupling matrix as it is directly included in the computation. This works well if the algorithm is applied in-situ and a coupling matrix taken when the radar was still calibrated can be provided. However, if the algorithm is applied as an initial calibration, it may be difficult to obtain an accurate passive coupling model for that system. We can combine aspects of the lattice and subtractive phase alignment algorithms to prevent the introduction of error. The lattice method relies on the geometric relationship that equally spaced elements (with consideration for polarization) can be assumed to have the same passive coupling between them. Instead, the passive coupling matrix can be used to find elements for which this assumption can be made.

There are two major advantages to this approach. First, the geometric assumption made by the lattice method may not always hold true. For example, elements closer to the edge of the array will begin to experience edge effects that change the behavior of the fields between elements [10], [12], [17]. Accounting for this effect requires additional logic and/or application of statistical metrics. Depending on the accuracy required for the application, the passive coupling estimate can be generated to include effects that are complex to describe in code. Second, by using the passive coupling matrix as a guide, it may be possible to find adequate pairings that would have been difficult to describe geometrically, do not have an obvious geometric relationship, or follow patterns that are difficult to describe. Since the model is used as a guide, it is not necessary for the passive couplings to be a numerically accurate prediction of the calibrated couplings. The model only needs to encode the relative

relationship between elements. If two couplings would be approximately equal when the radar is in a calibrated state, they should also be approximately equal in the passive coupling representation. Polarization is already taken into account ironically by ignoring it. The indexing could be randomized and the processes would work out exactly the same. There is not a point in the algorithm that treats the V pole of a particular element as having any relation to the H pol radiator of the same element. They are both just radiators with a given index, the system does not know which radiators go together. The only reason we structure H and V pol indexing is so that humans can interpret the data.

The implementation of the algorithm can be done in a variety of ways; for this research, it is implemented as follows. A transmit element and receive element are chosen to act as references. The algorithm then loops through every combination of pairings that involve the chosen transmitter, a common receive element, and another transmitter. If the phase of the passive coupling between the reference element to the receiver is close enough to the phase from the second element to the same receiver, eq. (2.1) is applied to align the transmitter to the reference. If the mutual coupling matrix is not updated each time, it is necessary to multiply by the alignment weight of the reference (solved in a previous loop). The newly aligned element is then added to the list of solved transmit elements to avoid solving it again. Once the loop is complete, every transmitter which can form a common receiver appropriately will be aligned to the reference for that loop. Then, a new reference is chosen from the list of solved elements. The same process is applied in reverse to solve receive weights. This process loops until all elements are solved.

The results of the application of this algorithm to a dataset collected on an

8x8 phased array system are shown in Fig. 2.9. This is an improvement to the results of lattice alignment on the same dataset shown in Fig. 2.2 because the guided alignment does not use the approximations and geometric assumptions. Instead, the assumption being made is that the relationships encoded by the passive couplings align with the relationships that exist in the real data.

2.6 Conclusion of Phase Alignment Discussion

The phase alignment techniques discussed in this chapter serve as a crucial pre-alignment step to prepare the system for more sophisticated calibration routines. As illustrated in Fig. 2.9, the results of the guided phase alignment method demonstrate that the phases of the elements are aligned within an acceptable range, ensuring that phase wrapping is avoided. However, it is important to emphasize that this pre-alignment step is not the final calibration. The goal of pre-alignment is to bring the system into a state where the phase variations are constrained, allowing the subsequent calibration algorithms to operate without ambiguity. While the pre-alignment process ensures that the system is prepared for further calibration, achieving precise alignment in both magnitude and phase requires more advanced methods. In the following chapters, we will explore existing calibration techniques and propose new approaches to address the limitations of current methods, particularly in handling incomplete datasets and avoiding reliance on geometric assumptions.

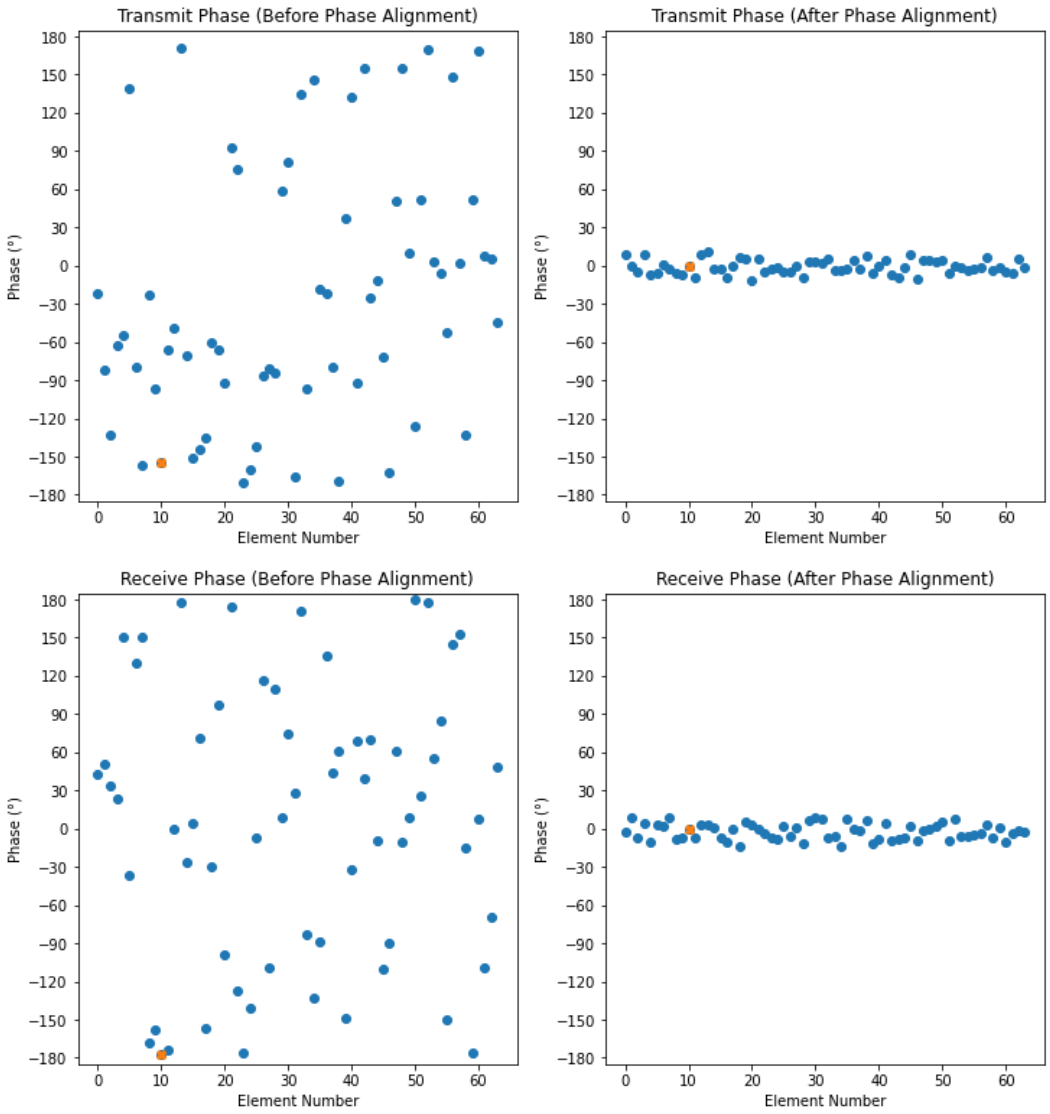

Figure 2.9: The result of guided phase alignment applied to a dataset gathered on an 8x8 phased array system. The element used as the first reference is marked in orange.

Chapter 3

Existing calibration Methods

3.1 Introduction

Calibration is essential for phased array radar systems to ensure precise control of signal phase and amplitude. Traditional methods, such as Park and Probe and Holographic Back-projection, rely on expensive and large equipment, making them costly and impractical for many systems. Mutual coupling-based calibration methods have emerged as a promising alternative, leveraging measurements between array elements to compute alignment weights without external equipment. However, these methods often depend on geometric assumptions and hard-coded logic, limiting their applicability to non-standard configurations and making them sensitive to incomplete data. This chapter explores the strengths and limitations of existing calibration methods. By analyzing these methods, we identify opportunities for developing robust, geometry-agnostic mutual calibration techniques. These techniques aim to overcome traditional limitations by leveraging passive coupling guidance matrices, providing flexible solutions for a wide range of phased array systems.

3.2 Constraints Limiting Existing Methods

There are various real-world conditions that make the implementation of these algorithms more difficult than they appear in simulation. Real world collected datasets are never fully completed. The specific hardware that enables the signal generation, transmission and reception will always impose restrictions on the dataset which must be accounted for to create a robust calibration algorithm. These constraints remove information from the system. Whenever the algorithm encounters a missing piece of information, the construction of the system of equations must be rerouted in some way to obtain that information elsewhere. Some of the effects that have been observed as part of this research are described in the following subsections.

3.2.1 Adjacent Element Censoring

The closer radiators are together, the stronger their coupling. In practice, this increased signal strength can saturate the receiver Low Noise Amplifier (LNA) [19]. In order to prevent the inaccuracies of saturation from distorting the later calculations or even making the system unsolvable, the collection process clips the measurement whenever the data is inside the saturated range. The dataset will have an unusable value for that particular index. This poses a problem if for example a [1,0] coupling set is used for the construction of the system of equations. One potential solution is to manually adjust the transmit and receive gain such that the system is operating within the linear region for a particular pairing. However, the next constraint limits the extent to which this linearization can be applied.

3.2.2 Range Limitation

Radiators that are farther apart will have a smaller coupling between them. If the radiators do not couple significantly, the mutual coupling measurement will generally be too small or may not be detected by the matched filter at all, resulting in a NaN value as shown with whitespace in Fig. 3.1. While adjacent element censoring can be solved by decreasing transmit and receive gain, those adjustments will also limit the range of each radiator, causing all pairings beyond a certain distance to be missing from the resulting dataset. There is a balance that must be achieved in order to properly linearize the system before taking measurements in order to have enough data [19].

3.2.3 Sub-array Censoring

Most of the datasets shown here were taken on the Horus Polarimetric Phased Array Radar [1], [20], [21] which has independent "Octoblades" that are the modular sub-unit of the front-end system. A single Octoblade has 16 ports supporting 8 dual-pole channels, with the primary drawback that these channels are not fully isolated. There is an existing problem that when taking measurements; there will be cross-talk between two channels on the same Octoblade. For this reason, the pairings between two radiators on the same Octoblade are censored in software. There should be future efforts to resolve this issue, but as of this writing, it is a constraint that must be taken into account. For experimentation purposes, multiple mutual coupling measurements can be taken on different ports.

3.2.4 Failed Elements, Missing Panels, and Masking

There is also the possibility of entire elements or sections of elements being faulty, dead, or disabled. It is common for systems to have modular panels of which some might not be installed. Algorithms must be capable of detecting these situations and avoiding those elements in the computation. It is useful to have the ability to apply a mask to the calibration operation such that masked elements are not included. With all of these ideas in mind, the next chapter proposes the Agnostic Mutual Coupling Algorithm and a direct method of including these functionalities will be shown. However, faulty elements may not be as simple to detect. There are many approaches to fault detection, but a straightforward approach is to use statistical thresholds which are based on previous observations of the system [12], [22]. This method does not work past a certain number of faulty elements because the rows and columns will be too far corrupted to get any meaningful data.

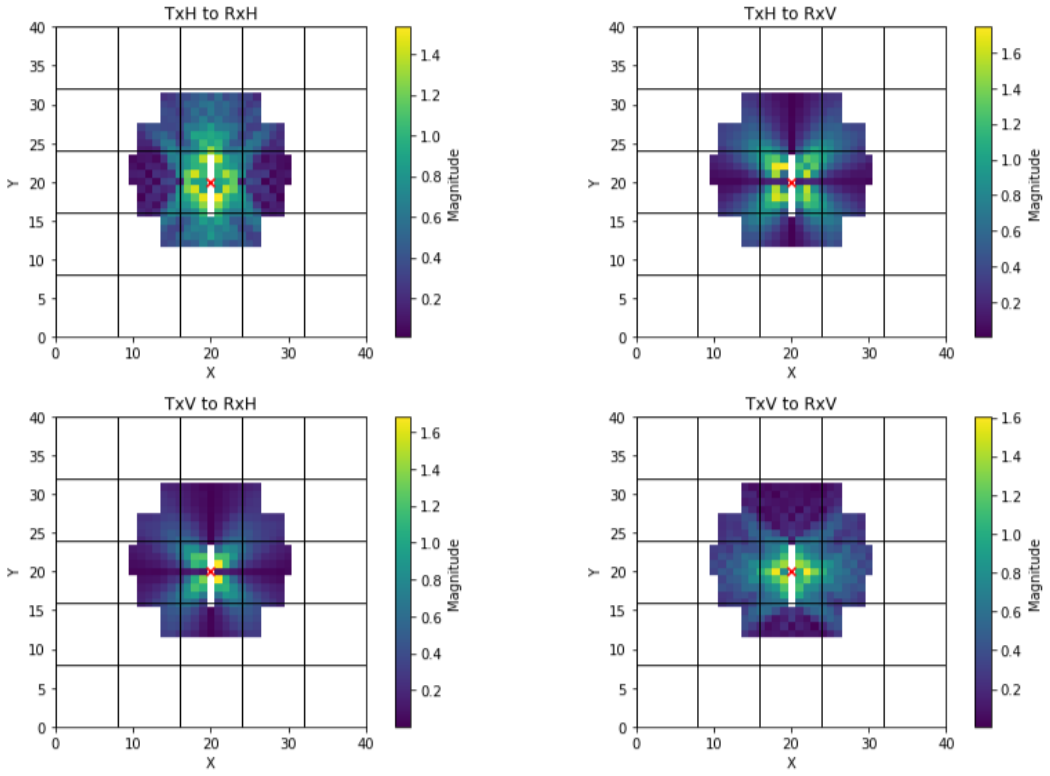

Figure 3.1: Plots of the magnitudes of the couplings for a 40x40 dual-polarized radar system. The couplings shown are all of the couplings that involve the center transmitter (marked with a red x). We can observe the effects of adjacent element saturation, range limitation, and Octoblade censoring. The full antenna is composed of modular 8x8 subpanels, which are indicated by black grid lines.

3.3 Park and Probe

Park and probe is a form of radar calibration that involves taking near-field measurements of a radar system using a reference antenna. In the simplest form, this can be a simple horn antenna placed a two to five wavelengths away or in the periphery of the radar panel [11], [23], [24]. The more common form of park and probe involves using a near-field scanner to move the reference antenna in a grid pattern above each radiator, capturing a set of near-field measurements to characterize each one independently [5], [11], [12], [19]. The system then evaluates the collected data to determine alignment weights. This is not feasible for many radar systems as it requires a near field scanner which is large enough to scan the full size of the antenna. These scanners are often sophisticated pieces of equipment that are particularly expensive. The park and probe method is also solely an initial calibration algorithm performed before deployment and cannot be applied as an in-situ method of calibration. Fig. 3.2 shows an example of a near field scanner that might be used for park and probe. These devices can be very large for certain applications, even up to the size of an airplane hanger or larger (must be larger than the antenna).

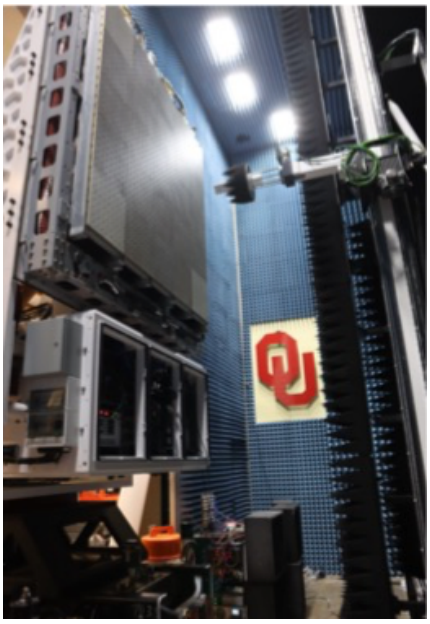

Figure 3.2: An example of a near-field scanner at the University of Oklahoma Advanced Radar Research Center (ARRC). This particular scanner was used to collect the park and probe measurements required to provide a ground truth for evaluating alignment weights.

3.4 Holographic Back-projection

Holographic back-projection is a near-field measurement technique used to estimate the excitation distribution of an array by reconstructing the radiated field at the array aperture. This method captures both the magnitude and phase of the near-field, allowing engineers to assess element-level variations in excitation and identify sources of distortion [5], [25]–[27]. It back-projects measured near-field data to the array plane, effectively reconstructing the aperture's excitation distribution. This allows for a more detailed assessment of the array's behavior while accounting for mutual coupling, scattering, and structural interactions that would be difficult to observe using traditional methods.

For effective application, the measurement scan plane must be precisely aligned with the array aperture, and deviations from planarity should be minimal to avoid phase errors. Additionally, an adequately sized measurement area is required to capture a sufficient portion of the radiated field and minimize truncation effects [27]. The method is particularly useful for phased array calibration, as it provides direct insight into element-level excitation variations, which can be iteratively adjusted to improve performance. Holographic back-projection has advantages over traditional calibration techniques, as it enables rapid near-field data acquisition over multiple beams and frequencies in a single scan. In addition, holographic back-projection uses the full aperture and captures the effects of temperature, reflections, scattering, and other environmental context that the traditional park and probe method does not characterize [27].

However, challenges remain, including sensitivity to environmental conditions such as temperature variations and the need for stable reference excitations during measurement [27]. This can be difficult to achieve as the array sizes get larger and the potential for large temperature fluctuations across the array become more likely. Furthermore, for power-amplified elements operating near saturation, the nonlinearity of amplifiers may necessitate iterative steps to achieve calibration [27]. This process requires a significant amount of expensive external equipment along with strict experimental constraints and somewhat intense processing. While this method provides a highly detailed characterization of array excitation, its complexity, equipment requirements, and experimental constraints make it impractical for many applications where less involved calibration techniques suffice.

3.5 Drone Calibration

In recent years, quad-copter drones have become very common for both commercial and recreational use. This availability makes them particularly accessible, especially for radar systems, which typically operate with large-scale budgets. In comparison to most RF measurement technologies, mounting an antenna on a quad-copter, along with some front-end electronics for data capture, is relatively inexpensive. As a result, there has been significant activity in the radar industry involving utilizing drones to obtain far-field data for calibrating radar systems. This approach is important to mention, as it does not suffer from many of the drawbacks discussed in this paper. Drones can be deployed for both in-situ and initial calibration, as they can take direct measurements that would often require expensive equipment, such as anechoic chambers equipped with field scanners. While there are still challenges associated with this approach, it strongly rivals the methods discussed here, particularly due to its extreme flexibility and low cost. Drones can even be multifunctional; many radar system tests can be performed using them. For example, to conduct object detection tests and reflectivity measurements, a metal sphere can be mounted to the drone and flown within the radar's field of view [28]–[31].

3.6 Embedded Feed Line Calibration

Embedded feed line calibration leverages integrated couplers and on-chip receiver circuits within phased array systems to enable precise, real-time calibration without external equipment. By embedding low-loss couplers at the input of each antenna channel, a known test signal can be injected directly into the array, allowing the system to measure phase and gain responses internally. This method uses built-in I/Q receivers to capture the response of each channel, facilitating accurate calibration of amplitude, phase, and frequency characteristics. Since the calibration occurs within the array itself, it compensates for variations due to temperature, aging, or environmental changes, ensuring consistent performance. This approach reduces the complexity and cost of external testing setups while enabling rapid, in-situ recalibration, critical for maintaining the accuracy of high-frequency phased arrays in radar and satellite communication applications.

This technique falls under the broader category of Built-In Self-Test (BIST) methodologies, where the system autonomously monitors and corrects its own performance. Embedded feed line calibration not only simplifies the maintenance process but also enhances system reliability by continuously tracking internal discrepancies and adjusting for any drift or degradation over time. This self-sustaining calibration mechanism is particularly valuable in environments where manual intervention is impractical, such as in spaceborne or remote sensing systems, where maintaining optimal antenna performance is essential for mission success [9], [11], [32]. The drawback of this method is the increase in hardware complexity and cost.

3.7 Coupling Set Calibration

Coupling set calibration is a method that leverages the mutual coupling measurements between array elements to compute alignment weights without the need for external equipment. This approach relies on the assumption that certain pairings of radiators have approximately equal passive coupling values due to their geometric similarity. By grouping these pairings into "coupling sets," the system can construct a solvable set of equations to determine the transmit and receive gains of each element as shown in [4], [6]–[8], [10], [12]. The following subsections detail the mathematical model, the process of creating coupling sets, and the statistical filtering techniques used to ensure the accuracy of the calibration. This method offers a promising alternative to traditional calibration techniques, particularly for systems where external equipment is impractical or costly. However, it is important to note that the accuracy of this method depends heavily on the validity of the geometric assumptions and the quality of the mutual coupling measurements.

3.7.1 Couplings Model

The uncalibrated coupling matrix, D, is given as an input which can be used to construct K defined as the ratio of two couplings as shown in eq. (3.1) [4], [6]–[8], [10], [12].

$$K_{C,D}^{A,B} = \frac{D_{A,B}}{D_{C,D}} = \frac{T_A S_{B,A} R_B}{T_C S_{D,C} R_D} \tag{3.1}$$

However, there are too many unknown variables present to solve a system comprised of these values. We must constrain these equations to a space that is solvable. We do this by limiting ourselves to only conditions where $S_{B,A} \approx S_{D,C}$. By constraining the system, the system becomes a lower dimensional vector space where it is feasible to construct a determined system of equations. This constraint adds new information to the system. The resulting expression is shown in equation (3.2) [6], [7].

$$S_{B,A} \approx S_{D,C} \implies K_{C,D}^{A,B} = \frac{D_{A,B}}{D_{C,D}} = \frac{T_A R_B}{T_C R_D} \tag{3.2}$$

At this stage, the system is formatted in a product expression. This is incompatible with traditional linear algebra methods, thus we must transform the system to an additive expression that can be mathematically expressed with matrix notation. To do this, we take the logarithm as shown in equations (3.3a) and (3.3b). The logarithm can be of any base; some implementations utilize base 2 as it provides the most dynamic range and least compression [10]. This results in an expression structured as linear system of equations.

$$log(\frac{D_{A,B}}{D_{C,D}}) = log(\frac{T_A R_B}{T_C R_D}) \tag{3.3a}$$

$$\implies log(D_{A,B}) - log(D_{C,D}) = log(T_A) + log(R_B) - log(T_C) - log(R_D) \tag{3.3b}$$

Using this expression, matrix notation can be applied by placing the *log* of each gain into a vector with a length equal to twice the number of radiators notated $\vec{\mathbf{E}}$. Along with this, we create a matrix composed of {-1,0,1}

which represents which of the gains are included in this particular equation and whether that gain was in the numerator or denominator of the original expression. Each row of this A matrix refers to a particular value of $\vec{\mathbf{K}}$ which we have included. The dimensions of $\bar{\mathbf{A}}$ is the number of equations added to the system by twice the number of radiators. To complete the expression, we place the corresponding value of K into a vector with length equal to the number of equations in our system notated $\vec{\mathbf{K}}$. The result is the expression shown in eq. (3.4). The goal of calibration is to solve for the unknown circuit gains of each transmitter and receiver chain referred to the antenna terminals, thus the expression is solved by applying the Moore-Penrose pseudo-inverse of $\bar{\mathbf{A}}$. See Appendix A for a description of the Moore-Penrose pseudo-inverse.

$$\vec{\mathbf{K}} = \bar{\mathbf{A}} \cdot \vec{\mathbf{E}} \iff \vec{\mathbf{E}} = \bar{\mathbf{A}}^{+}\vec{\mathbf{K}} \tag{3.4}$$

These coupling measurements on their own are not enough to create a fully determined system of equations as the system would always be circularly defined. To constrain the solution space, we will often provide at least one transmit and one receive gain. In order to factor this into the system of equations, we could simply create a row of $\bar{\mathbf{A}}$ with a single 1 placed in the appropriate column and then place the *log* of the known gain in the corresponding row of $\vec{\mathbf{K}}$. For this implementation, the columns of $\bar{\mathbf{A}}$ corresponding to the known gain(s) are placed into a separate matrix, $\bar{\mathbf{B}}$. The columns corresponding to the unknown gains are then placed in a new matrix $\bar{\mathbf{G}}$. The log(s) of the known gain(s) are placed in a separate vector notated $\vec{\mathbf{E_2}}$ and likewise the unknown gains will be represented with $\vec{\mathbf{E_1}}$. The resulting expression is shown in eq. (3.5) [6], [7], [10], [12].

$$K = \bar{G} \cdot \vec{E_1} + \bar{B} \cdot \vec{E_2} \iff \vec{E_1} = G^{+} \cdot (\vec{K} - \bar{B} \cdot \vec{E_2}) \tag{3.5}$$

$$\begin{bmatrix} 1 & -1 & 0 & 0 & 0 & -1 \end{bmatrix} \begin{bmatrix} log(T_A) \\ log(T_C) \\ log(T_D) \\ log(R_A) \\ log(R_C) \\ log(R_D) \end{bmatrix} + \begin{bmatrix} 0 & 1 \end{bmatrix} \begin{bmatrix} log(T_B) \\ log(R_B) \end{bmatrix} = \begin{bmatrix} log(\frac{D_{A,B}}{D_{C,D}}) \end{bmatrix} \tag{3.6}$$

3.7.2 Coupling Sets

One method of ensuring the constraint of always using pairings that have equal passive couplings is always met is to create coupling sets [4], [10], [12]. These are groups of geometrically similar pairings that are likely to have equal passive couplings because of geometric rules. For example, a pairing of the H-pol of element (0,0) to the H-pol of element (0,2) is likely to have the same passive coupling as a pairing of the H-pol of element (5,0) to the H-pol of element (5,2). It is important to note there is a dependence on polarization, which means that simple Euclidean distance cannot be used. Therefore, consideration must be applied to the polarization of transmit and receive [13]. A depiction of coupling sets is shown in Fig. 3.3. Each pairing within a given set is added to the system of equations together with each of the other pairings in that set. The major advantage of this method is that systems of equations can be built according to simple indexing rules. However, there are a few major

drawbacks. First, this method is likely missing pairings that would have had equal passive couplings which did not follow the simplistic geometric pattern. Second, indexing logic is required to implement this method which is highly sensitive to missing information. If the algorithm encounters values that do not exist in the measurement, additional logic must be implemented to handle every possible situation. When bundled with the necessity for an algorithm that can perform on variable-size arrays, the logic just to manage proper indexing becomes very intricate. Third, the pairings within a given set are often not actually equal to one another. For example, pairings closer to the edge of a radar panel have edge effects that make the passive coupling between them different than a similar pairing closer to the center [4], [10], [12], [22]. An inspection of the coupling values encompassed by a given set will reveal that the assumption of equal passive couplings is often not true. This problem can however be compensated to achieve good results by using statistical rejection methods discussed in a later section. Finally, this method is heavily dependent on having an equally-spaced, planar phased array system. Systems with more complex geometry require more complex methods of creating coupling sets or require different algorithms entirely.

3.7.3 Coupling Set Results

A 4x4 panel example will be shown to measure the performance of the coupling set method. A few different coupling sets are grouped together using arrays of pairings. A series of *for* loops are then used to add these pairings into the system of equations. In order to understand the minimum possible error using this method, a system with unity passive couplings will be used. This checks that the system of equations are fully determined and solvable. Any error in the more realistic simulation is the result of inadequately chosen pairings and not an indeterminate system. These results will then be compared to results that were computed using more realistic passive coupling values obtained through simulation as shown in Table 3.1. In the next section, we will discuss a statistical method to improve this error.

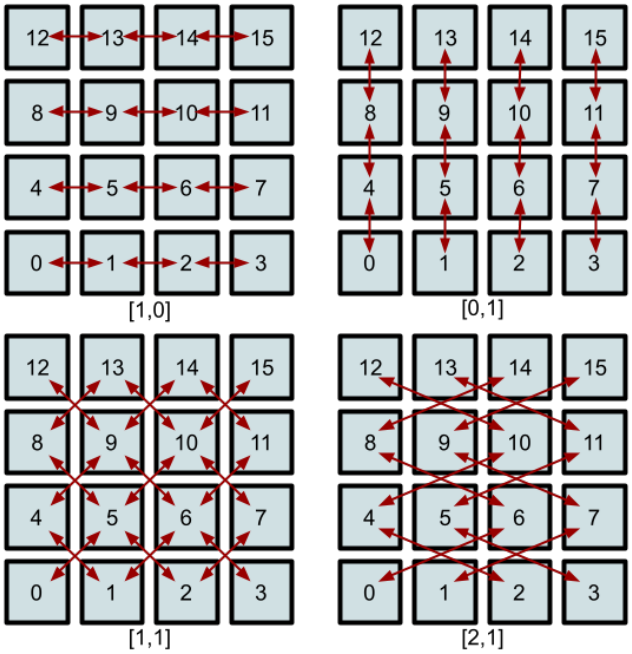

Figure 3.3: An illustration explaining the notation for coupling sets. The coupling set notation takes the form of [horizontal, vertical] using translation distances. Coupling that have the same geometric spacing are in the same set.

Unity Passive Couplings	**Magnitude (%)**	**Angle (deg)**
Average Error	1.71e-13	3.21e-14
Max Error	4.71e-13	1.59e-13
Realistic Passive Couplings		
Average Error	34.98	9.21
Max Error	70.21	21.39

Table 3.1: Error metrics for coupling set calibration, one with unity passive couplings and the other with more realistic couplings. From this information, we can see that the error is essentially zero indicating that if the coupling set assumptions are correct, the results should have no error.

3.7.4 Statistical Set Filtering

When evaluating a particular coupling set, it is often the case that the pairings contained within are not as close as is initially assumed. One method of compensating for this ambiguity is by creating a function that will collect the couplings that would have been used by that set and then apply statistical metrics to determine if the set falls within acceptable bounds [4]. If the set fails to meet the statistical conditions, the pairings included in that set will not be included in the system of equations. This is not a catch-all method, but has been found to provide good results when testing on equally spaced planar systems. Many metrics can potentially be used but for the purposes of this research, it was found that standard deviation, minimum, and maximum are sufficient for achieving accurate calibration. The exact thresholds for each of these will depend on the context of the dataset, but a standard deviation of 2, minimum coupling ratio of 1e-30, and maximum coupling ratio of 1e-10 are generally good starting points. These numbers may be highly subject to variation across different radar systems due to a variety of factors that are beyond the scope of this paper. It is recommended to find thresholds that are tailored to the datasets that are generally output by the system in question. If the metrics are too restrictive, the resulting system will be under-determined as there were not enough included pairings to obtain an adequate solution. Conversely, if the metrics are not restrictive enough, the system will be determined but will provide inaccurate results as the assumption of equal passive couplings was too far from reality.

3.8 "Reference Based" Mutual Coupling Calibration

The method often used for mutual coupling calibration in practice as of this writing is generally referred to as "Reference Based Mutual Coupling Calibration" in conversation. This name is not ideal as almost every mutual coupling calibration algorithm involves some kind of reference and this is not even the only algorithm that involves a reference matrix. It is a variation of the process described by [8], [11], [12], [33], in particular [34]. Şeker simply refers to the "Mutual Coupling Method" which would now be too broad a term. There does not appear to be a more precise term for this early implementation in the literature. The algorithm applies eq. (3.5). Instead of using a measured mutual coupling matrix, the measurement matrix is divided by a carefully selected reference matrix that is usually a coupling measurement taken while the radar was in a calibrated state. This simplifies the mathematics as there is no need to find pairings that would have equivalent passive couplings or a need to divide by another coupling to cancel passive coupling terms at all. In addition, the algorithm is implemented on a per-system basis with factors such as the $\bar{\mathbf{A}}$ matrix being mostly pre-defined. There are some implementations that apply a graph theory Python library to check if the missing information hinders the solvability of the system, but this is not the same structured algorithmic creation of a system of equations that is described here. This method is very geometry-dependent and will only apply to similarly spaced planar phased arrays. The deployment environment also must be electromagnetically similar to the environment that the reference matrix was collected. All of the mutual coupling methods described here are extensions of the concepts learned from these implementations.

3.9 Coupling Linearization

In order for the operations used throughout this document to be applicable, the linearity of the system must be ensured. If the transmit power of every element is increased by 10 dB, there should be a resulting 10 dB increase for every coupling. In this thesis, two methods of achieving this linearization will be discussed. First, digital predistortion which involves altering the outgoing signal in amplitude and phase before it reaches the amplifier such that the amplifier will distort it to the desired signal. Second, a more manual form of correction which involves sweeping across amplifier settings and evaluating the coupling data with each setting to find the optimal configuration to perform coupling measurements. Without these linearization steps, non-linear effects can occur. For example, if the transmit magnitude is too low, the accuracy of the received signals will drop, this can be observed from the drop in signal to noise ratio (SNR) that occurs with lower transmit power. This will result in psuedo-random increases in coupling magnitude with respect to increase in transmit power. If transmit power is too high, the amplifier will be pushed into the saturation region. This will result in a smaller increase to coupling magnitude than the transmit power was increased by. Both of these effects can be seen in Figs. 3.4 and 3.5. If a linear change to gain magnitude does not result in a linear change to the coupling, then modeling the gains as simple multiplication of a complex value is no longer a useful prediction of reality. All of the derivations throughout this document rely on that linear assumption, thus it must be ensured before the algorithms can be applied.

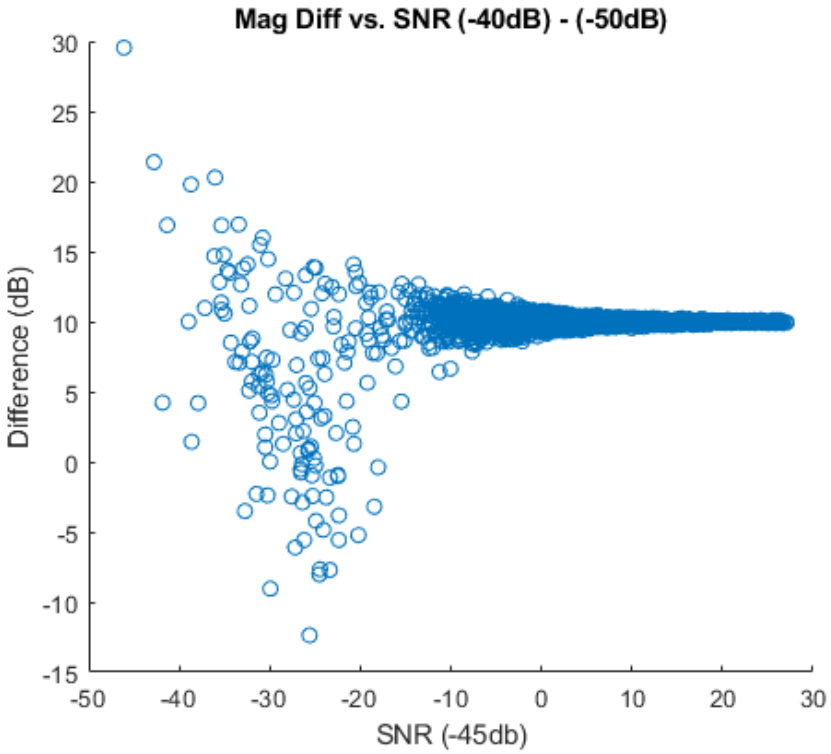

Figure 3.4: The division of the coupling dataset taken at a transmit setting of -40 dBFS divided by the coupling dataset taken at -50 dBFS. An increase of 10 dB on each transmitter should ideally result in an increase of 10 dB to each coupling. In this experiment we can see that we do in fact observe this behavior for couplings which have an adequate signal to noise ratio. As the signal to noise ratio deteriorates, we no longer observe this linear relationship.

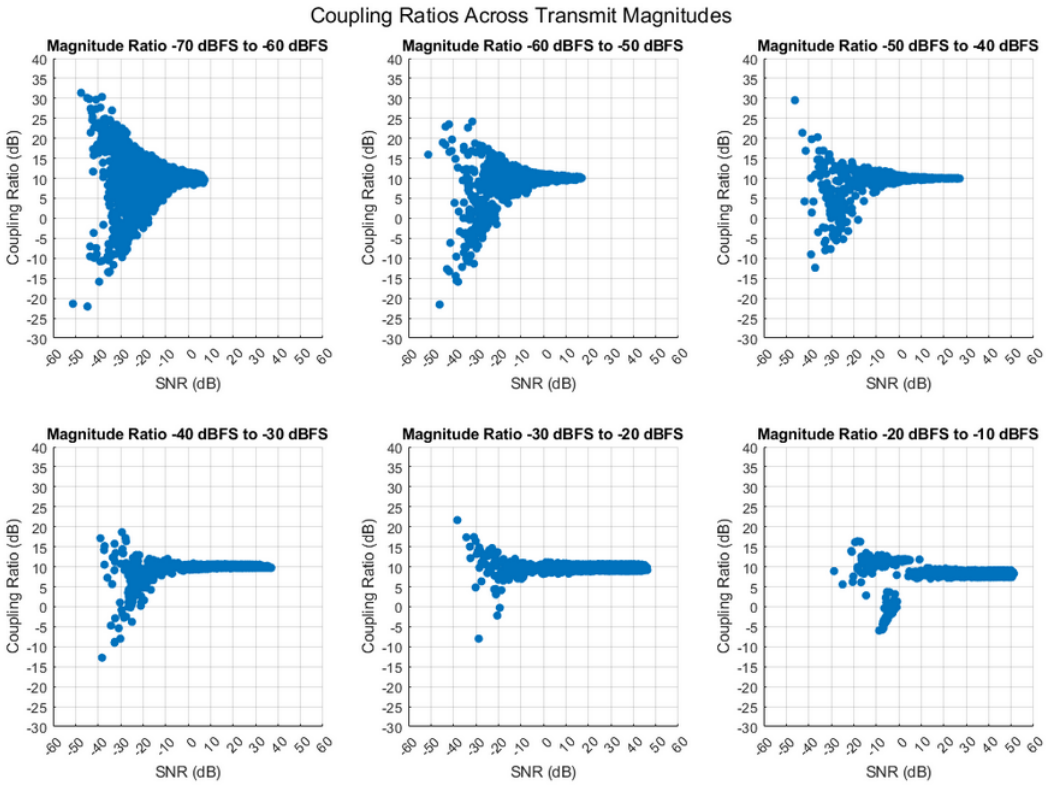

Figure 3.5: The comparison of coupling data across various ranges of transmit gains. When operating outside of the linear region it can be observed that the linearity becomes chaotic even for values which have high SNR. The non-linear effects of both ends can also be observed from these plots. At lower magnitudes, the SNR is further and further reduced such that more of the couplings are taken with low SNR. At the highest magnitude, the effect of amplifier saturation can be observed as even the high SNR couplings are consistently below the expected 10 dB.

3.9.1 Digital Predistortion

Digital Pre-Distortion (DPD) is a technique used to linearize power amplifiers in communication systems by compensating for the nonlinear distortions introduced by the amplifier. In particular, the memory polynomial model is a commonly used model that captures both the static nonlinearity and the dynamic memory effects present in amplifiers.

Memory Polynomial Model

The output of an amplifier operating in the saturation region, $y(n)$, can be expressed as a function of the input signal, $x(n)$, using a memory polynomial model. This model includes both nonlinear terms and memory effects, where the memory terms account for the dependence of the output on past inputs. The general form of the memory polynomial model is given by [14], [35]:

$$y(n) = \sum_{p=1}^{P} \sum_{m=0}^{M-1} h_{p,m} x(n-m)|x(n-m)|^{p-1} \tag{3.7}$$

where:

- $y(n)$ is the output of the power amplifier at time step n.
- $x(n)$ is the input signal to the power amplifier.
- $h_{p,m}$ are the complex coefficients of the memory polynomial model, which need to be estimated during the DPD training process.
- P is the maximum degree of nonlinearity.
- M is the memory depth or the number of past input values considered.

The memory polynomial model effectively captures both the nonlinear behavior (through the polynomial terms) and the memory effects (through the past input samples) of the amplifier. The term $x(n-m)|x(n-m)|^{p-1}$ represents the p-th order nonlinear interaction of the input signal at time $n-m$.

In DPD, we aim to design a pre-distorter function, $d(n)$, such that the input to the amplifier, $x(n)$, is preprocessed to cancel out the distortions introduced by the PA's nonlinearity. The pre-distorted signal is defined as:

$$x(n) = d(n)s(n) \tag{3.8}$$

where $s(n)$ is the desired linear output signal, and $d(n)$ is the digital predistorter applied to $s(n)$ to compensate for the amplifier's nonlinearity. By applying the memory polynomial model and optimizing the coefficients $h_{p,m}$, the predistorter ensures that the amplifier output $y(n)$ is approximately equal to the desired output $s(n)$:

$$y(n) \approx s(n) \tag{3.9}$$

Coefficient Estimation

To implement DPD, the coefficients $h_{p,m}$ of the memory polynomial model must be estimated. This is typically done by collecting training data, which includes the input signal $x(n)$ and the corresponding amplifier output $y(n)$. The coefficients are optimized by minimizing the error between the desired linear output $s(n)$ and the actual amplifier output $y(n)$. This can be done using least squares or other optimization techniques to solve the following minimization problem:

$$\min_{h_{p,m}} \sum_{n} |y(n) - s(n)|^2 \tag{3.10}$$

Once the coefficients are estimated, they can be applied to the pre-distorter to linearize the amplifier output. There are multiple potential ways to obtain these coefficients from a training dataset. For this discussion, we will focus on a method that utilizes linear algebra to compute the least squares solution using a Moore-Penrose pseudo-inverse as shown in detail in [14]. First, we must represent the memory polynomial model in terms of matrices as shown in eq. (3.11). Here N represents the length of the input vector and M represents the memory order. The "delay matrix" $\bar{X}$ shown in eq. (3.12) contains rows that correspond to the memory polynomial term associated with each coefficient. The columns of the matrix are across samples in time.

$$\vec{y} = \begin{bmatrix} y(M) \\ y(M+1) \\ \vdots \\ y(N-1) \\ y(N) \end{bmatrix} = \bar{X} \begin{bmatrix} h_{00} \\ h_{01} \\ \vdots \\ h_{0m} \\ h_{10} \\ \vdots \\ h_{k0} \\ h_{k1} \\ \vdots \\ h_{km} \end{bmatrix} \tag{3.11}$$

$$\bar{X} = \begin{bmatrix} x(M) & \cdots & x(1)|x(1)|^k \\ \vdots & & \vdots \\ x(N) & \cdots & x(N-M+1)|x(N-M+1)|^k \end{bmatrix} \tag{3.12}$$

Once $\bar{X}$ and $\vec{y}$ have been populated with training data, eq. (3.11) can be solved by multiplying the Moore-Penrose psuedo-inverse $\bar{X}^+$ with $\vec{y}$ to obtain the memory polynomial coefficients. Note that if multiple datasets are used, the same format can be applied as long as attention is paid to using the corresponding input vector for that dataset when populating the rows of $\bar{X}$. This method may also become infeasible for applications with significantly large datasets as the inverse becomes much more difficult to compute in a reasonable time. For characterizing amplifier distortion, the amount of data needed to train the pattern is small enough to apply this method. This methodology is importantly applicable in both directions. In order to model the behavior of an amplifier, the $x(n)$ corresponds to the input of the amplifier. However, in order to create a digital pre-distortion model, the datasets are actually applied in reverse. The function of the pre-distorter is to apply a change to the signal such that when the amplifier distorts the signal, it returns the signal to the original desired signal. The DPD model can be thought of as the inverse of the amplifier distortion. There are some extra normalization constraints which are discussed in more detail in [14]. These conditions are beyond the scope of this paper, but any code and data shown will be utilizing those considerations.

State Matrix Representation

As used in [36], one of the simpler ways to obtain model predictions after the memory polynomial coefficients are found is to use a state matrix as shown in eq. (3.13) and to represent the coefficients using the coefficient matrix, $\bar{H}$, which organizes the coefficients as a 2D matrix instead of the linearized form used before. The memory polynomial model is by definition not time-invariant. Therefore, to compute the output of the model, some information about the prior states of the model must be known. The state matrix representation makes this process very straightforward as it maintains all of this contextual information without needing to store the history of the signal and recompute the terms used in previous calculations for every model prediction. Each column of the state matrix corresponds to a particular sample of the dataset. Each time a new piece of data is fed into the model, the state matrix shifts left, dropping the oldest vector, and appends the state vector for the current time step on the right. The coefficient matrix represents information learned during training of the model and does not change after training [1]. The results of applying this model to data collected from an amplifier is shown in Fig. 3.6.

$$y(n) = \bar{H} \cdot \bar{S}_n \tag{3.13}$$

$$where,\ \bar{S}_n[k,m] = x(n-m)|x(n-m)|^k \tag{3.14}$$

[1] Some applications of memory polynomial may be implemented so that the model can train itself as it makes predictions. This is useful for designing "online learning models".

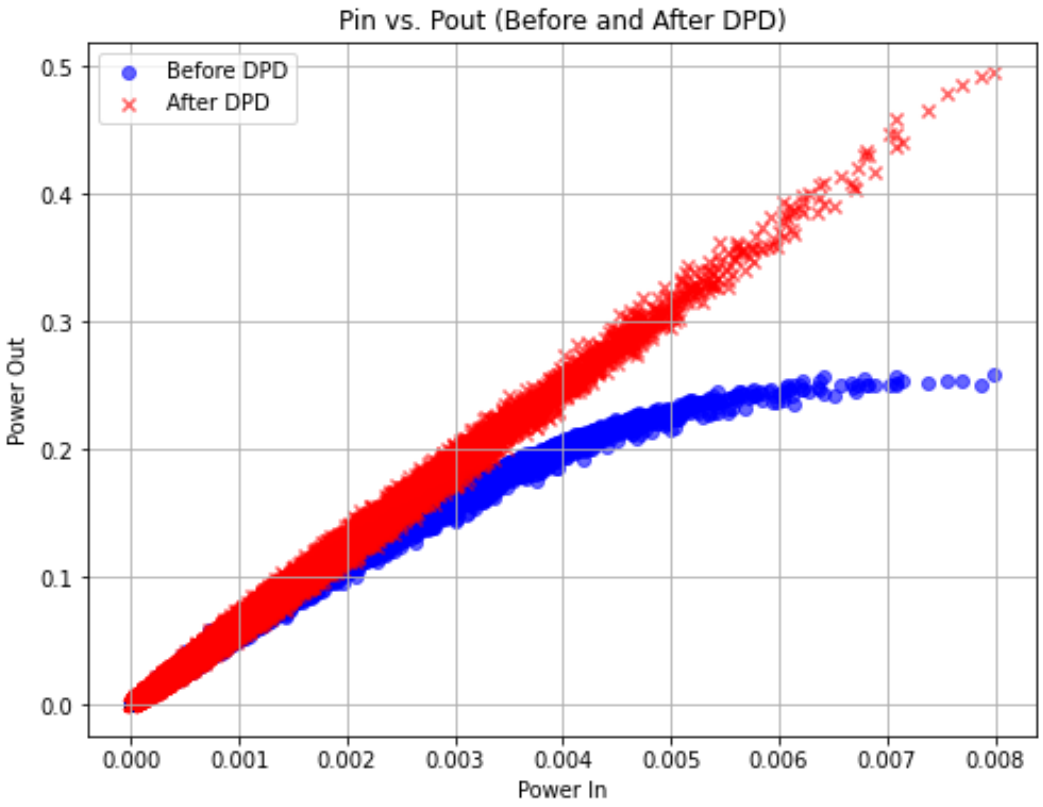

Figure 3.6: The relationship between the input power and the output power of data taken from a power amplifier. The blue circles represent the amplifier's output power before applying DPD, while the red crosses represent the output power after applying DPD. The application of DPD reduces the nonlinear distortion, as evidenced by the improved alignment with the ideal linear response. This demonstrates the effectiveness of the model in linearizing the amplifier's behavior.

3.9.2 Sweep Linearization

For many applications, it may not be feasible to implement digital pre-distortion. A quick and simple method to obtain linearized coupling data is to sweep across various gain magnitudes to find a state that is linear. In many systems the receive gains are purely digital, in which case this sweep becomes simpler to perform as there is only one optimization dimension. By comparing coupling matrices taken at different gain settings, the linearity of the couplings can be measured by dividing one matrix by the other. If every transmitter increases by 10 dB, every coupling should increase by 10 dB. This gives a metric to test if a given range of gains are operating in the linear region. In the system on which this methodology was tested, the appropriate linear range was determined to be between -30 dBFS and -50 dBFS, as shown in Fig. 3.4. There is a direct correlation with the signal-to-noise ratio which must be taken into account for calibration. As can be seen in Fig. 3.5, low SNR couplings do not exhibit linear behavior. Calibration algorithms should filter out couplings which have low SNR measurements if the system provides SNR measurements.

Phase variation must also be taken into account. An increase in magnitude should not affect the phase of coupling measurements. To ensure this stability condition, the angular matrix of one dataset can be subtracted from the other to obtain the plots shown in Fig. 3.7 and Fig. 3.8. In the case of the datasets collected on the particular system shown, the phase difference was determined to be within tolerance for high SNR measurements. The same pattern observed with magnitude comparisons can be shown here, that low SNR measurements violate the stability conditions.

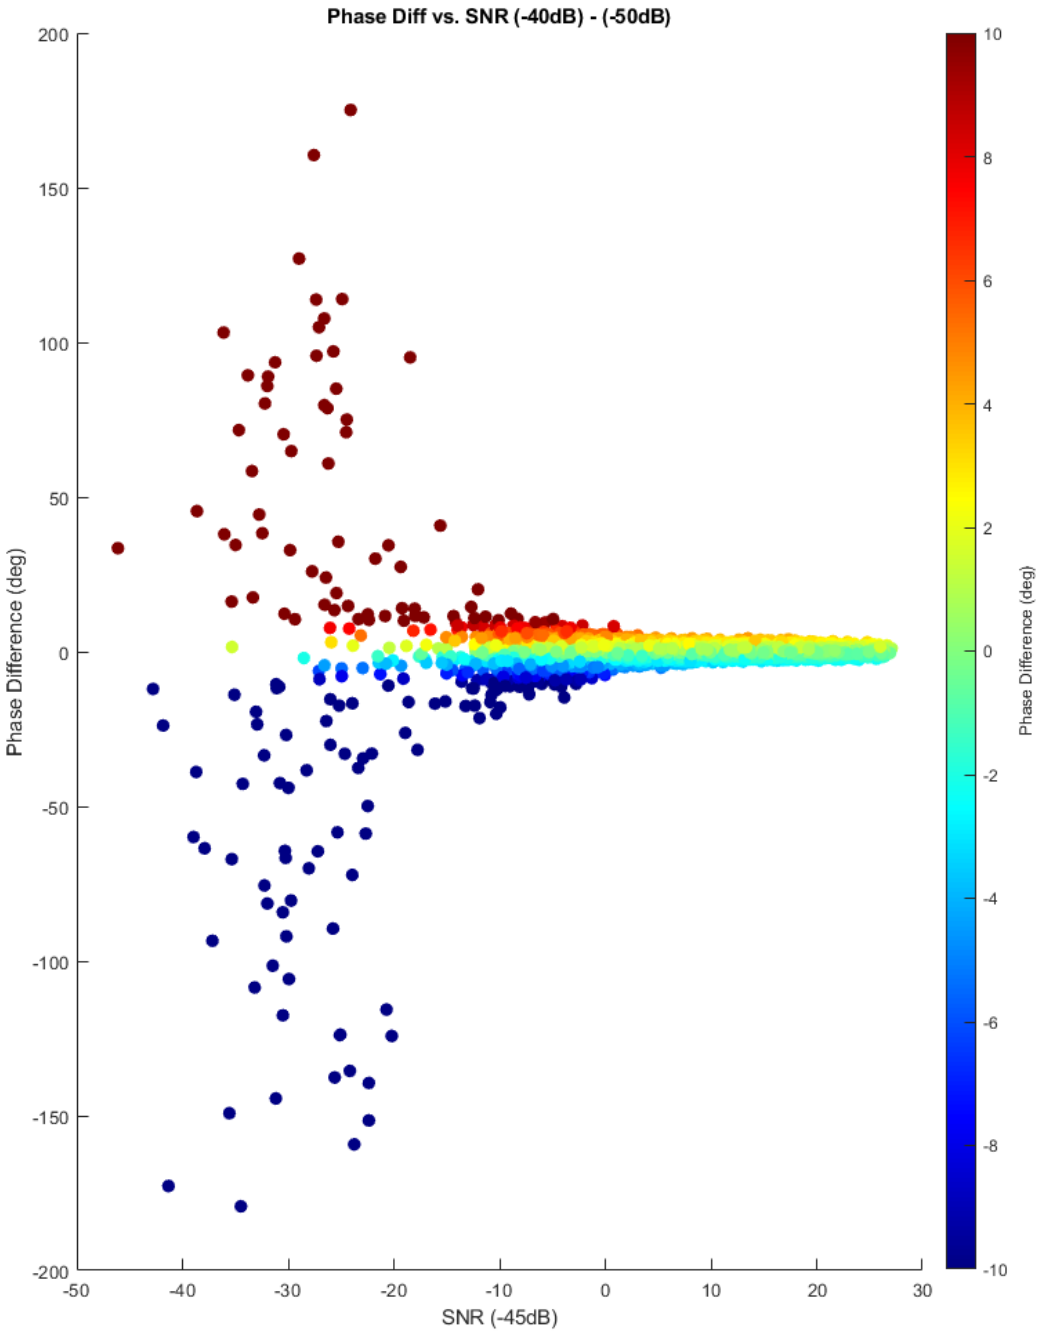

Figure 3.7: The difference in phase between coupling data taken when the radar was set to a transmit gain of -40 dBFS compared to -50 dBFS. As can be seen in the plot, the change in gain magnitude had minimal effect on the phase of the couplings as long as the SNR of the particular measurement was high enough.

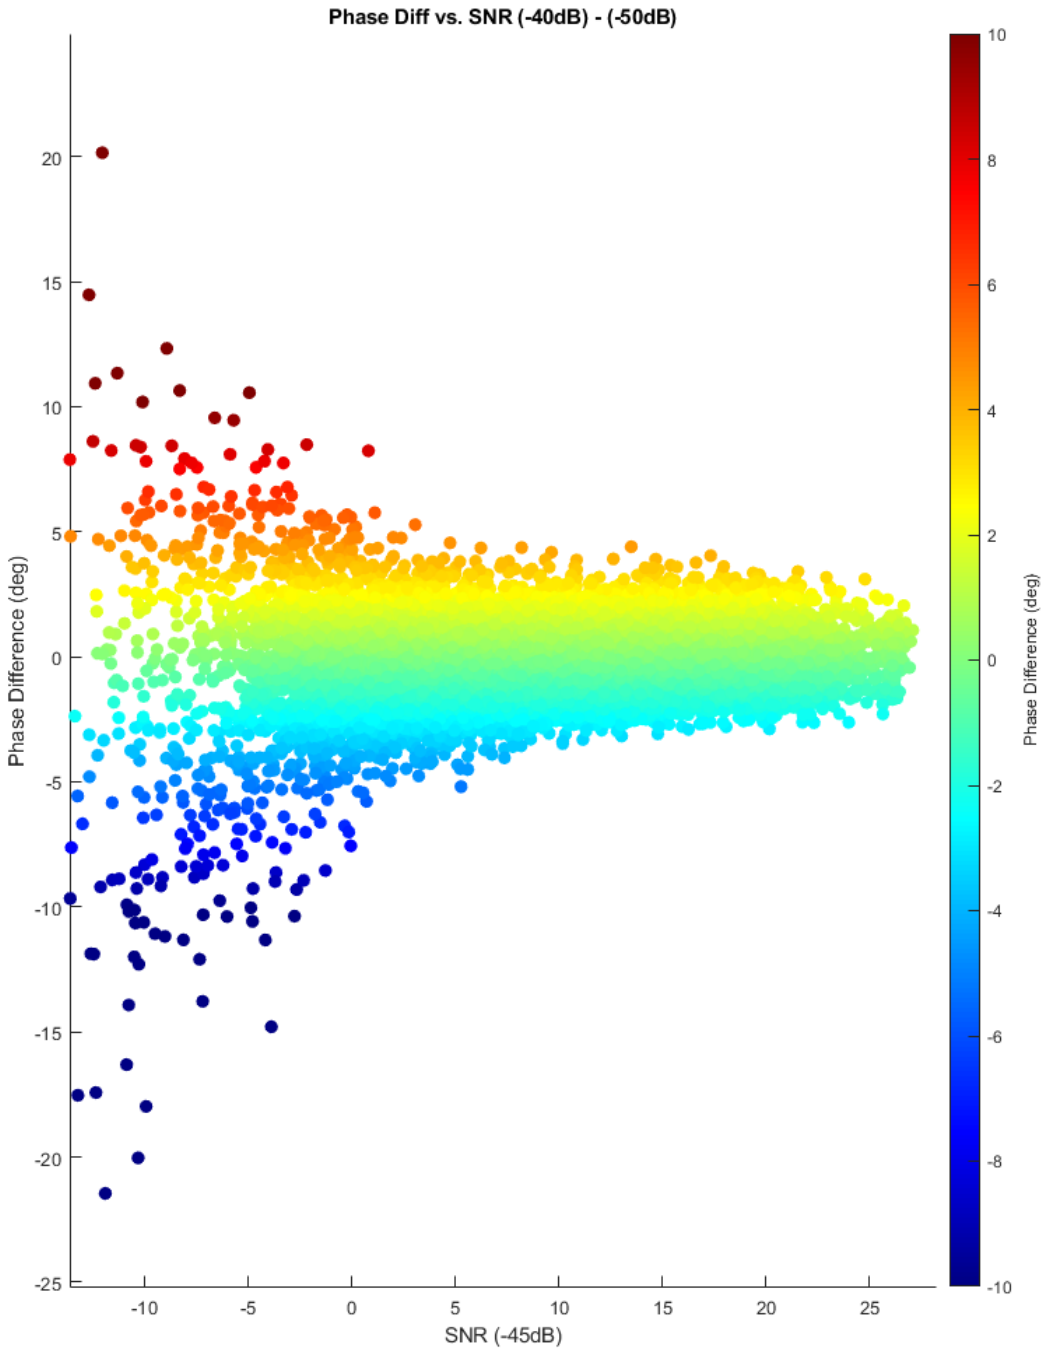

Figure 3.8: The difference in phase between coupling data taken when the radar was set to a transmit gain of -40 dBFS compared to -50 dBFS. As can be seen in the plot, the change in gain magnitude had minimal effect on the phase of the couplings as long as the SNR of the particular measurement was high enough. Here the plot has been zoomed in to show the high SNR measurements.

Chapter 4

Geometry Agnostic Mutual Coupling Algorithm

4.1 Adding Couplings to the System

To create a framework for managing the system of equations, a Python class is created. Within this class, the attributes of the particular system which are necessary to achieve calibration are defined. For example, the number of elements or the $\bar{\mathbf{A}}$ matrix itself. There is not enough information within the existing matrix to know which pairings were used in every case. The pairings that have been used so far are stored as metadata so that when it comes time to update the $\vec{\mathbf{K}}$ vector with new values utilizing the same system of equations, there is no need to reconstruct the matrix. An object variable to track the number of times that each gain is used in the system of equations is also implemented. This can give the operator a sense of how well distributed the system of equations is. The term "element" is used to imply a position of the array, distinct from the term radiator. For dual pol systems, each element has two radiators (H and V).

Next, a function must be implemented to add new couplings to the system. Given the constraints previously mentioned, it is efficient to apply filtering functions at this level, ensuring that higher-level processes don't need to handle checks like whether the pairings involve NaN values. Additionally, we can apply other constraints based on the provided data, such as ensuring that passive coupling values are sufficiently close to each other. As we will see later, applying these filtering metrics at this stage simplifies the implementation of algorithms.

The function returns 1 if the pairing is successfully added to the system, and 0 if any filtering stage rejects the pairing. One major advantage of this approach is the separation between the passive coupling (calibrated) matrix estimate and the actual computation. The passive couplings are never directly included in the computation but serve as a guide for constructing the system of equations. The accuracy of the passive couplings is not critical, as long as pairings that are close in the real world are also close in the passive coupling dataset.

It is important to note that the subtractive and lattice phase alignment methods are sensitive to the accuracy of the passive coupling phase. This is because unlike the process used here, the angular information of the passive coupling matrix is directly included in that calculation.

4.2 Brute Force Method

The most inclusive algorithm is to brute force every possible combination of radiators and pass them into the object. This has the benefit of including every viable pairing into the system of equations, but for larger systems will take a considerable amount of computation time ($O(N^4)$) and RAM in order to generate $\bar{\mathbf{A}}$. It is possible to simply save the calibration object to a file after the system of equations has been constructed, then update $\vec{\mathbf{K}}$ with fresh data when it comes time to apply a calibration. This requires running an entire brute force process, but it can be done beforehand instead of having to generate everything in real time. In this case, the pairings used in the pre-generated system must be present in the new dataset.

4.2.1 Bruteforce Results

The Bruteforce method can be evaluated in the same way that the coupling set method was evaluated in Table 3.1 as shown in Table 4.1. First, a control is obtained by evaluating a system where the mutual coupling measurement is simulated with passive couplings that are set to one. This forces all pairings in the system of equations to be equal, demonstrating what accuracy the system of equations can achieve if all pairings that were assumed to be approximately equal really were equal. Note that the ones matrix is only for the simulation of the measured (uncalibrated) mutual coupling matrix, the same passive coupling estimate matrix is used for both experiments so that the resulting system of equations is the same. It is observed in this case that there is negligible error in the case of Unity Couplings, indicating that the system is determined and any error in the realistic passive coupling results is

due to improper choice of pairings. When we evaluate the error associated with realistic passive couplings, we still observe a non-negligible error as with the coupling set data, however these error margins are within more acceptable tolerances. Additionally, this error can be reduced by setting the angle and magnitude thresholds to lower values (until the system is so exclusive that it is under-determined). An example of the gains at each step is shown in Fig. 4.1 and the state of the mutual coupling matrix at each step in Fig. 4.2.

Unity Passive Couplings	**Magnitude (%)**	**Angle (deg)**
Average Error	2.36e-13	3.78e-14
Max Error	5.23e-13	1.13e-13
Realistic Passive Couplings		
Average Error	3.44	0.84
Max Error	6.96	1.66

Table 4.1: Error metrics for the brute force method, one with unity passive couplings and the other with more realistic couplings. From this information, we can see that the error is essentially zero indicating that if the pairing assumptions are correct, the results should have no error. The next entry has the same results but uses more realistic passive coupling values. A non-negligible error is still observed, however, the errors are within an acceptable tolerance. This error can be improved through more restrictive thresholds.

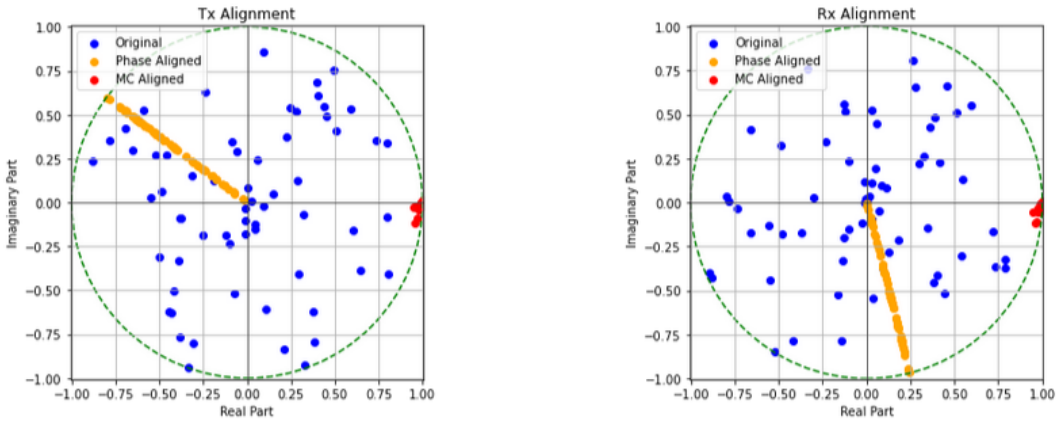

Figure 4.1: A scatter plot depicting the gains at each stage of the calibration process. Initially, the radar is uncalibrated with randomized gains, indicated in blue. After phase alignment, each gain is aligned to the reference element, shown in yellow. Finally, a brute-force alignment is applied, resulting in the gains being approximately one shown in red. Note that the Tx and Rx angles are different as the receivers are aligning to the Rx reference and the transmitters are aligning to the Tx reference.

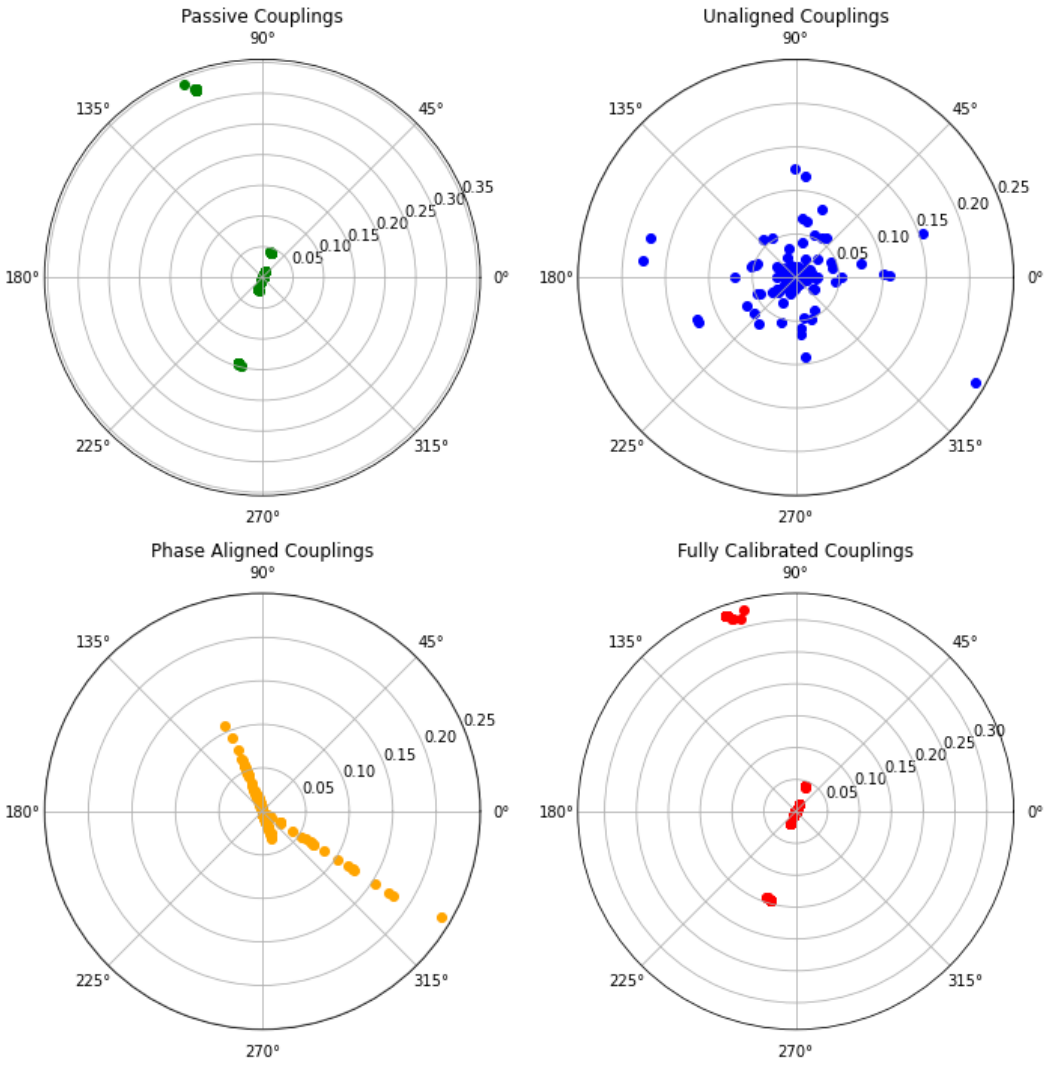

Figure 4.2: A scatter plot depicting the state of the mutual coupling matrix at various states throughout the calibration process. The top-left plot depicts the passive coupling coefficients of the radar system. The top-right plot is the mutual coupling matrix when the system is uncalibrated. The bottom-left plot depicts the mutual coupling matrix after the array has been phase-aligned. The data points become linear and at roughly the appropriate phase. The bottom-right plot depicts the state of the mutual coupling matrix after the brute-force method has been applied. If the solution is correct, the mutual coupling matrix will be equal to the passive coupling matrix.

4.3 The Role of Reference Gains

As discussed throughout this paper, all mutual coupling-based algorithms require a known reference gain to achieve accurate calibration results. Without this reference, the resulting system of equations becomes circularly defined. The maximum rank that the system can attain without a reference gains is equal to the number of unknowns (alignment weights) minus two. However, there may be situations where an operator is unable to take this measurement. In such cases, obtaining a partial calibration is preferable to remaining entirely uncalibrated.

The simplest approach to obtain calibration results without a reference is to assume that the reference has not drifted significantly from its factory-calibrated state and to input its complex gain as one. This assumption causes the algorithm to output the alignment weights relative to the reference as shown in Fig. 4.4. Consequently, we can apply the known gain after computing the results by multiplying the obtained weights with the reference weight. This process is mathematically identical to the separation method of including the reference described in eq. (3.5). If these relative weights are used as alignment weights, the elements will be aligned to their respective reference, which may or may not be a calibrated state.

Alternatively, if no reference is applied at all and the under-determined system is solved anyway, the transmit and receive gains will still align with one another. However, they will not achieve a configuration that calibrates the system, as illustrated in Fig. 4.3. To improve the magnitude, we can compute a scaling factor. This factor is obtained by calculating the square root of the average ratio between the calibrated mutual coupling matrix and the partially

aligned mutual coupling matrix (i.e. the measured matrix divided by the relative alignment weights). This approach recognizes that the relative alignment weights will yield a subsequent mutual coupling matrix with a different overall scale than the calibrated matrix we aim to achieve. By re-scaling the alignment weights to match the magnitudes of the matrices, we can enhance the estimate.

The square root is taken because this scale difference must be corrected through the multiplication of transmit and receive gains, meaning that each gain must contribute to the scaling factor. While the magnitude alignment will still be imperfect in this scenario, it will be an improvement over the initial relative alignment weights. This improvement is illustrated in Fig. 4.3. The greater the magnitude difference between the transmit and receive gains prior to re-scaling, the further the resulting gains will deviate from a magnitude of one after the operation. There does not seem to be a method to similarly correct the misaligned phases without providing the phase of at least one reference element as the phases of the two matrices would be the same.

The resulting alignment state when applying this method is not actually calibrated. For instance, there will be a net phase offset applied to all data collected by the system according to the sum of the phases of the transmit and receive gains. In a calibrated state, this sum would be zero as any non-zero phase in transmit is countered with an opposite phase in receive. Additionally, if the transmit and receive did not converge to the same magnitude, there would be an overall magnitude shift applied to all data collected. While this is unacceptable for most applications, it is an improved state to the totally uncalibrated state that the system started in.

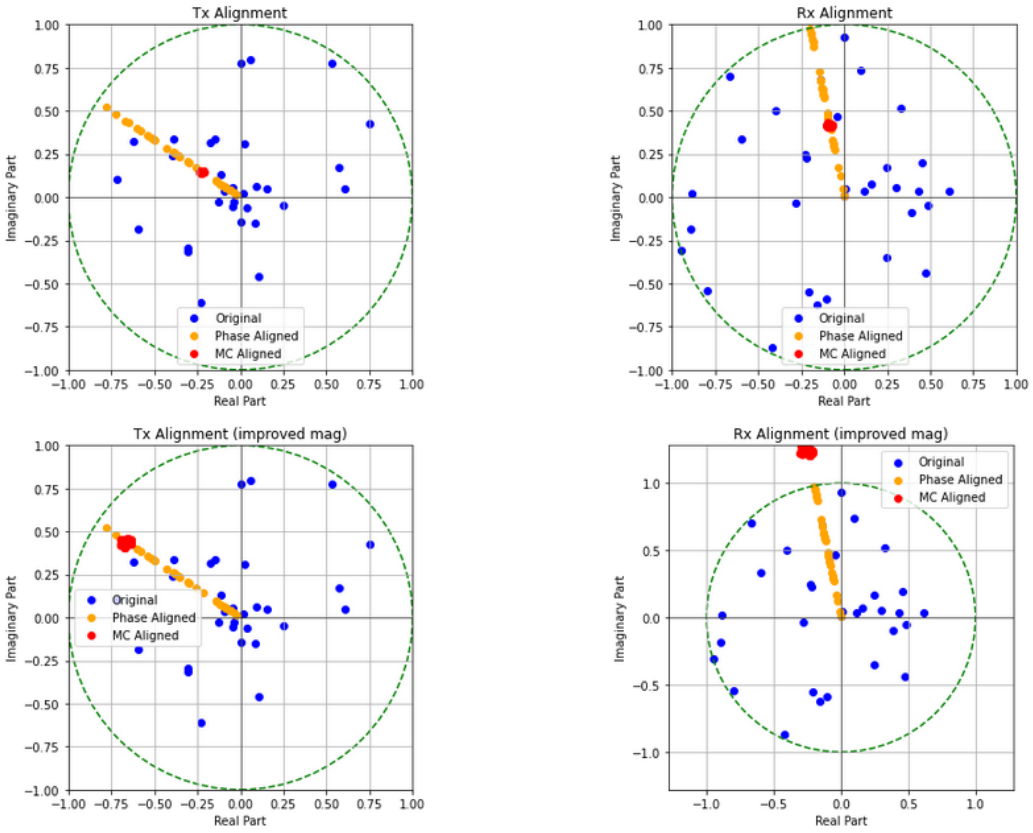

Figure 4.3: When no reference gain is provided, the system still aligns all of the transmitters to a particular value and all of the receivers to a particular value, but the converged values are not a calibrated state. The magnitude can be corrected by applying a scale factor which is equal to the resulting coupling matrix and the reference (calibrated) matrix.

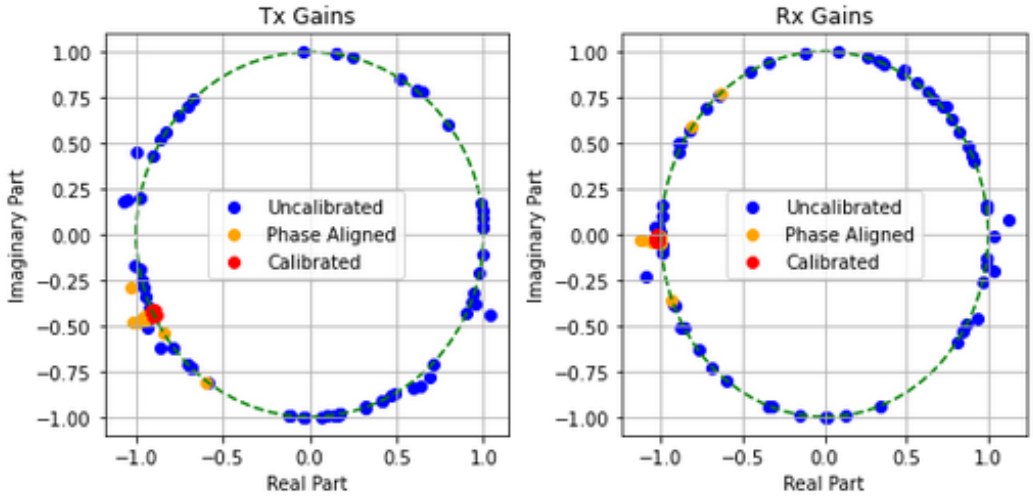

Figure 4.4: The result of calibration performed on data from a dual-pole 8x8 system if the reference gains are assumed to be one. All elements are aligned to the complex gain of the reference element. The alignment weights can then be divided by the reference gain to align all the elements to one.

4.4 Experimental Process & Results

To gather the real-world data that was used throughout this thesis, the following experiment was performed. First, as discussed in the section "Coupling Data Linearity", mutual coupling datasets were collected at various amplification settings with an 8x8 equally spaced phased array radar system. The data shown in Fig. 3.5 shows the results of that process. Through these plots we can observe the effects of saturation at the highest amplification; the 10dB increase resulted in less than 10dB due to the amplifier limits being reached. At lower amplification, we can observe the expected reduction in overall SNR which results in more non-linear couplings. The system was determined to be linear between -30 dBFS and -50 dBFS. In practice, this process should be automated as a preparation step to the calibration process.

After the system is placed into a linear state, the uncalibrated coupling data is collected. A near-field scanner was then used to perform a park and probe alignment of the radar. The near-field samples used to determine the appropriate alignment weights were saved to provide a ground truth for the algorithmic solution to obtain error metrics and confirm validity of results. After the system was successfully calibrated, another mutual coupling data collection was performed to be used as a passive coupling matrix. When using coupling data taken while the radar was in a calibrated state as the passive coupling, the algorithm is being applied as an in-situ calibration.

The weights obtained through park and probe can be applied to the uncalibrated matrix to verify that they provide an accurate linear mapping from the uncalibrated state to the calibrated one. The calibration weights can also be visually confirmed by comparing the uncalibrated couplings from a given transmitter as shown in Fig. 4.5 to the same data after the alignment weights are applied shown in Fig. 4.6. The uncalibrated phase plots are psuedo-random while the calibrated state has an apparent structure. This confirms the validity of the data for use as a ground truth metric. The results of brute-force applied to each possible phase alignment is shown in Figs. 4.8 through 4.10 to demonstrate that the phase alignment does not need to be a perfect alignment to solve effectively. It just needs to eliminate phase wrapping. Without providing a reference, the accuracy of the relative solution can also be verified by dividing the uncalibrated couplings by the matrix obtained through multiplication of calibrated couplings and the solved gains as shown in Fig. 4.7.

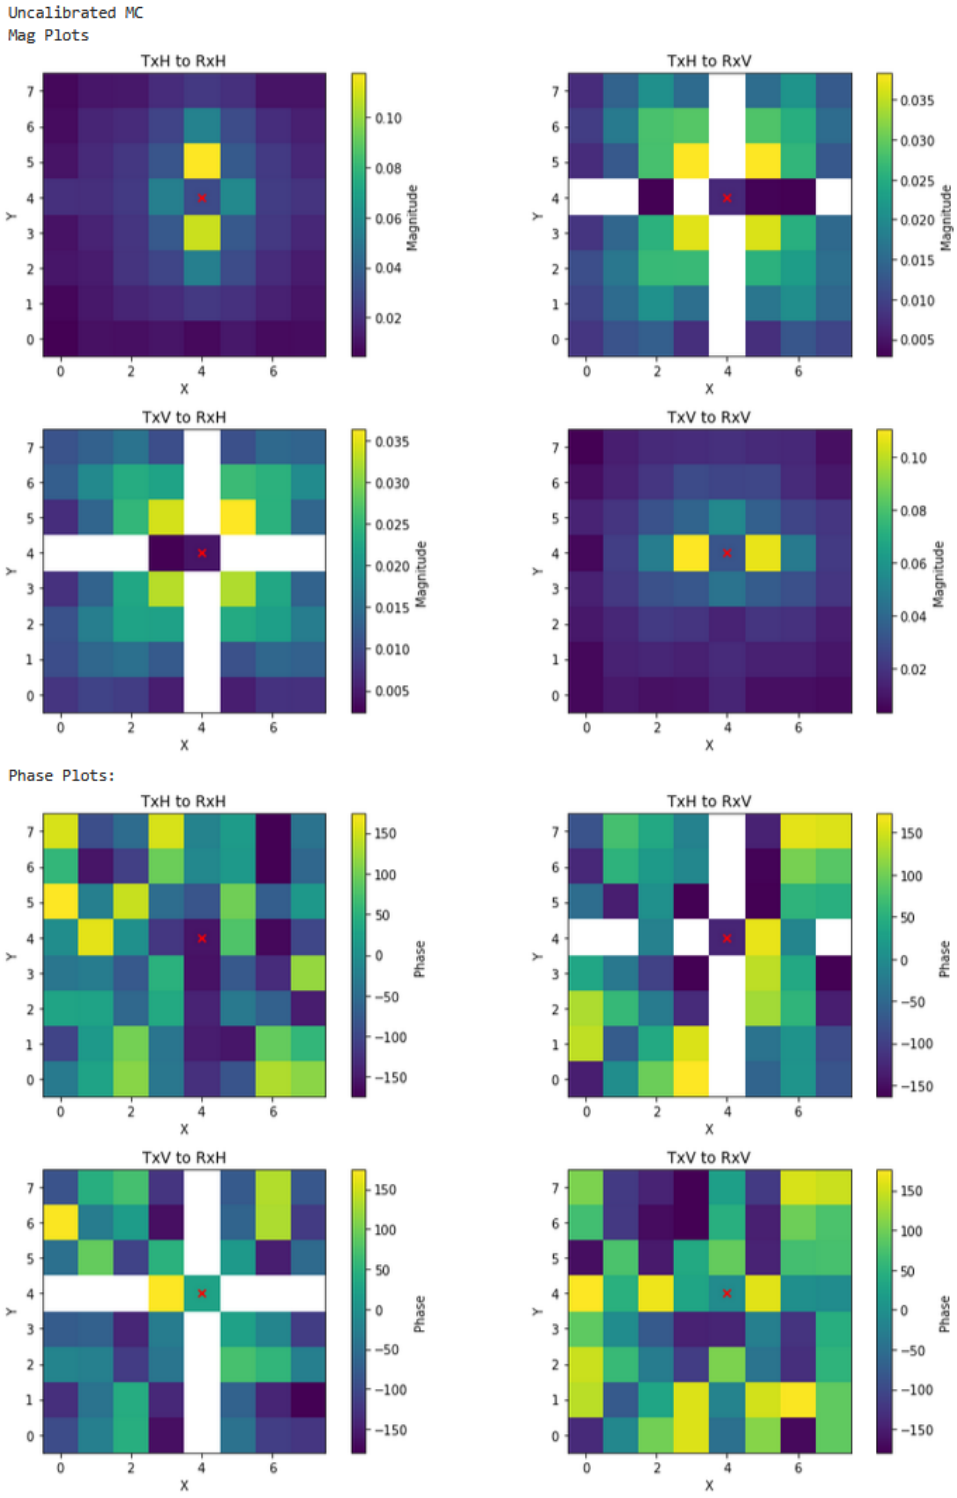

Figure 4.5: The couplings associated with transmit element (4,4) measured before calibrating the array. Low SNR measurements are masked.

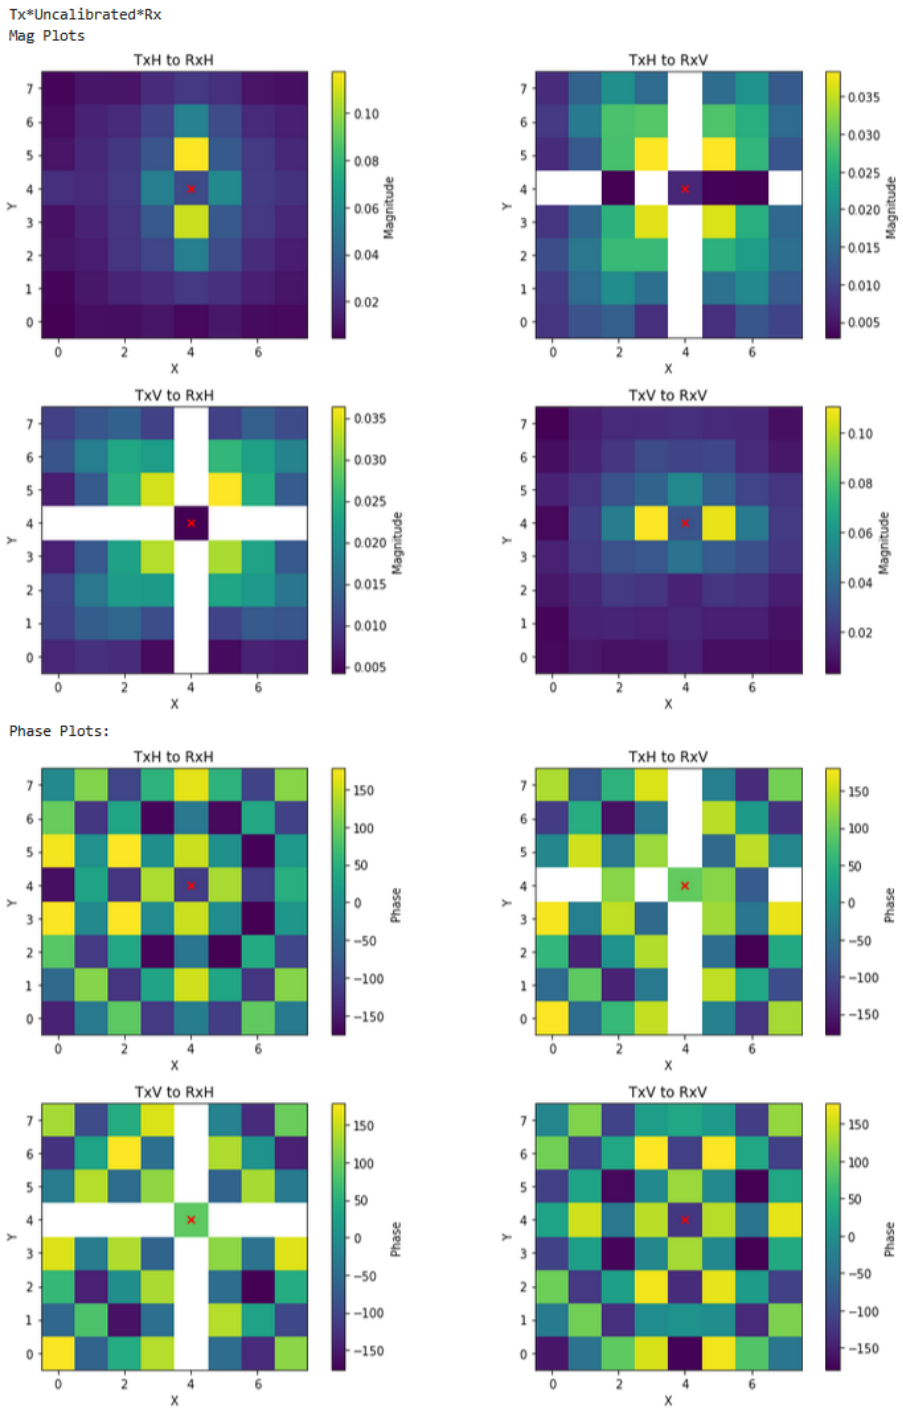

Figure 4.6: The result of multiplying the uncalibrated coupling data shown in Fig. 4.5 with the alignment weights produced through park and probe measurements. Low SNR measurements are masked.

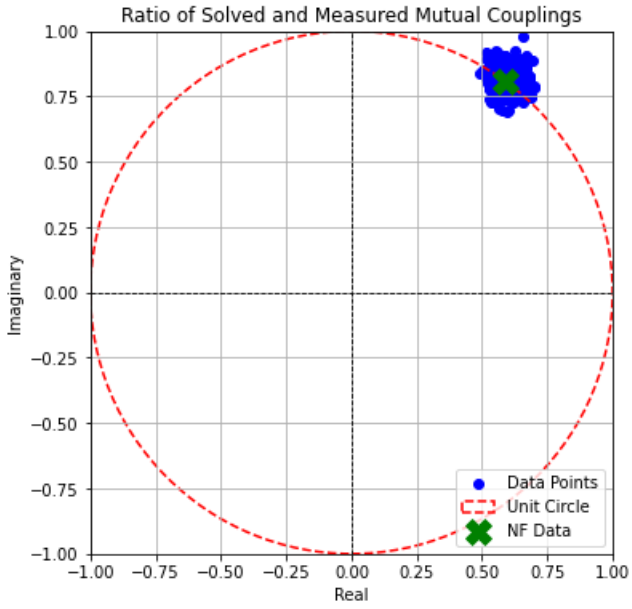

Figure 4.7: Solution accuracy can be verified by taking the calibrated matrix and multiplying by the solved gains to obtain a matrix that should be identical to the uncalibrated matrix. If the reference gains are set to one, the resulting matrix will be offset from the uncalibrated matrix by a constant factor equal to the multiplication of the transmit and receive references. The plot shows the division of the two matrices.

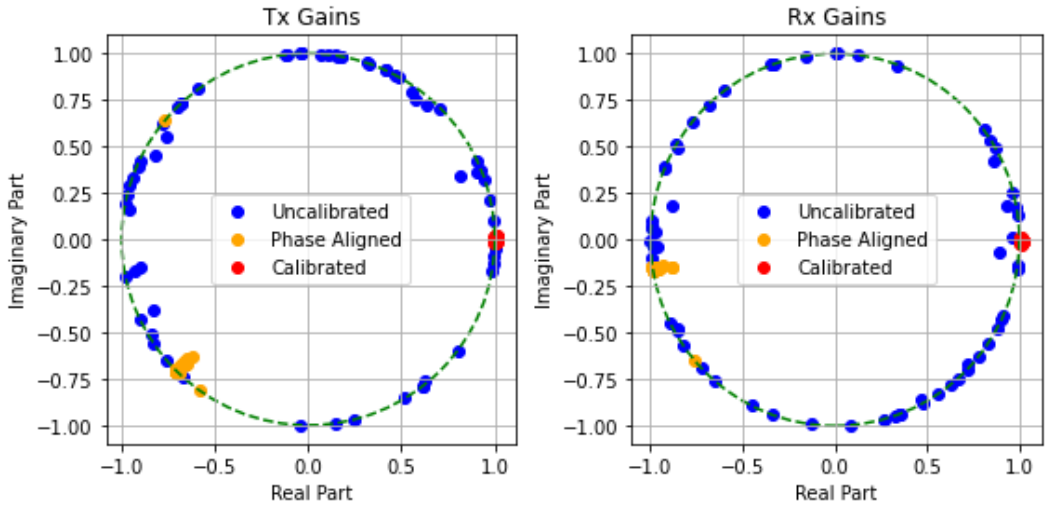

Figure 4.8: Blue: The ground truth gains obtained by inverting the alignment weights obtained through park and probe. Orange/Yellow: After subtractive phase alignment correction is applied. Red: After phase correction and alignment weights have been applied.

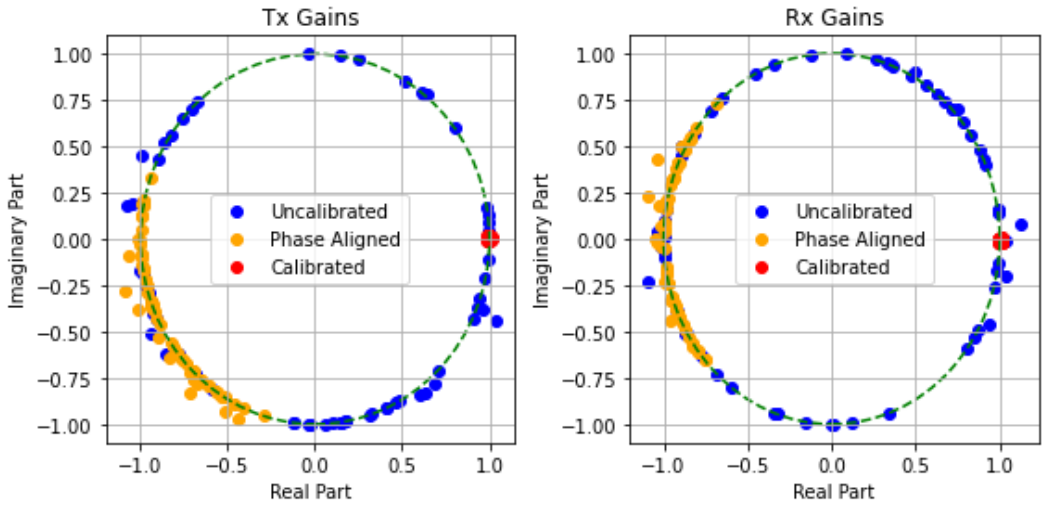

Figure 4.9: Blue: The ground truth gains obtained by inverting the alignment weights obtained through park and probe. Orange/Yellow: After lattice phase alignment correction is applied. Red: After phase correction and alignment weights have been applied.

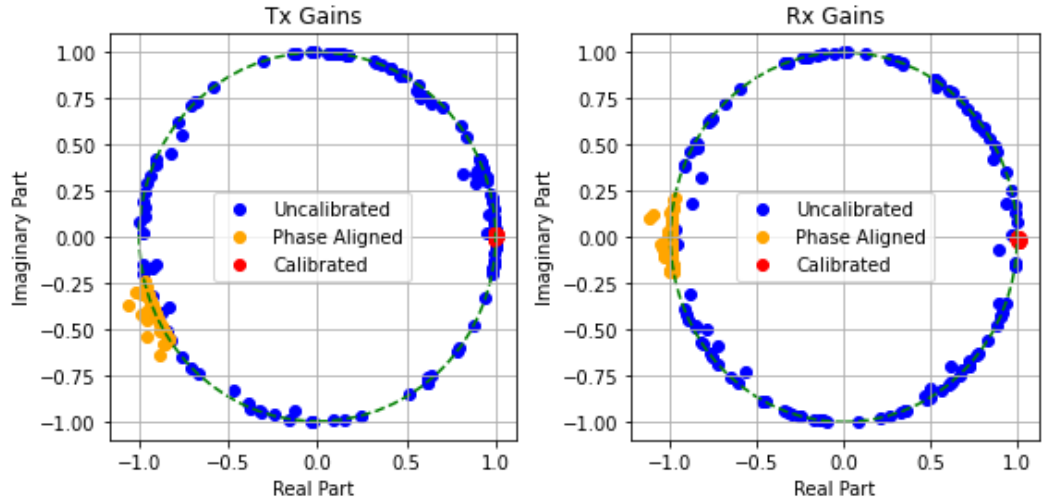

Figure 4.10: Blue: The ground truth gains obtained by inverting the alignment weights obtained through park and probe. Orange/Yellow: After guided phase alignment correction is applied. Red: After phase correction and alignment weights have been applied.

Chapter 5

Partitioning

5.1 Panel Subdivision Method

It is generally not feasible to perform most calibration algorithms on arbitrarily large arrays. As the dimensions of the array increase, the amount of RAM needed even to store the coupling matrices themselves becomes very large. It is therefore necessary to partition the computation into subunits, not only to limit resource usage, but also to enable parallelization of the computation. The most straightforward way to do this is by subdividing the array into panels. This is intuitive as large phased array radars are often designed with modular panels anyway.

First, a reference is selected for every panel, generally in the same location on every panel. Next, the alignment process is applied to all of the reference elements as if they were a single panel as shown in Fig. 5.1. For example, if we have an array made of 8x8 panels that are arranged as a 4x4, this reference array would be calibrated as if it were a 4x4 array. This aligns each reference to the main reference element. Then, each subdivision can be processed independently, aligning to its reference. This allows the system to be solved in parallel and does not require the entire coupling matrix to be in RAM.

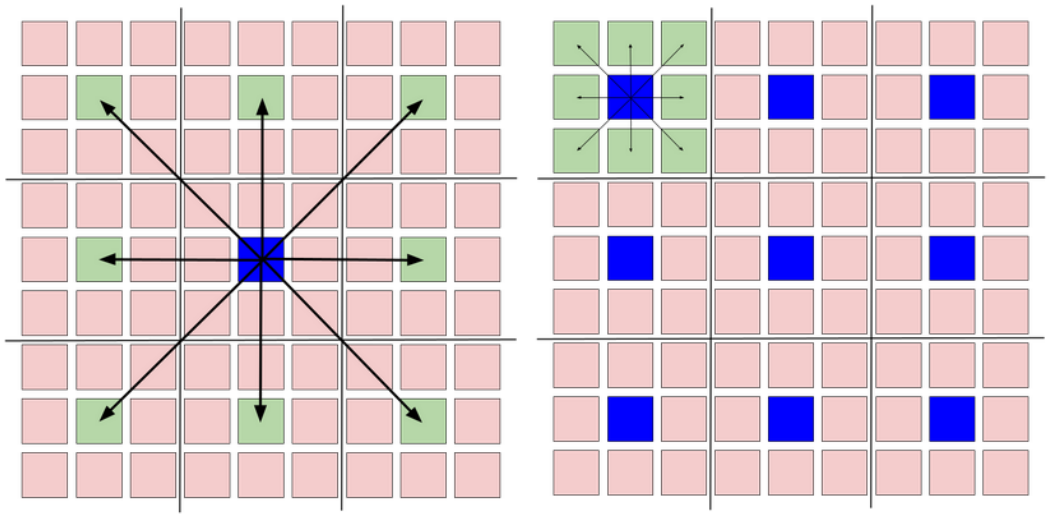

Figure 5.1: A depiction of the panel subdivision method. Blue represents elements with known gain, green represents elements that are being solved in the current iteration, and red represents elements that are yet to be solved. The array is subdivided into panels such that first the reference for each panel can be solved, then each panel can be solved independently.

The critical drawback of the panel subdivision method is that it is not feasible for range-limited datasets. This is because the initial reference alignment step cannot be completed since the references are generally far enough apart that there are no valid coupling values connecting them. It is possible that a higher transmit gain could be used for the reference alignment step, but higher amplification generally results in more significant non-linearity introduced by the amplifiers and the amplification used for the increased range would need to be factored into the results in order to obtain proper alignment. For very large systems, this amplification may even push the amplifiers into saturation. It is still possible to subdivide the array without the reference alignment step if a reference is provided for each subdivision.

5.1.1 Dynamic Loading

For very large arrays the RAM required to store the couplings and variables involved in solving the array can become quite large. For systems with limited RAM, it may be necessary to only load the partition currently being solved into RAM. The trade-off is that this will increase the number of read/write operations to the permanent storage device. As a consequence, dynamic loading is only cost beneficial if the permanent storage is a tolerant medium such as a solid-state drive (SSD) or flash memory. Hard disk drives will deteriorate rapidly from the increased frequency of reading and writing. The cost of increased wear to the hard disk will likely outweigh any savings from using less RAM. However, in some systems it may be infeasible to provide the necessary amount of RAM, in which case dynamic loading may enable the system to solve the array at all. This offloading to storage can be performed on a software level for greater control or by simply increasing the amount of storage dedicated to the swap space partition (virtual memory in Windows).

5.2 Agnostic Subdivision

While there is still research to be done in the direction of agnostic subdivision, an algorithm called connective calibration which produces accurate results for most elements will be described. First, the algorithm picks a number of references set by the function caller by determining the number of connections each element has. The elements with the most connections are selected to act as the references for that solve cycle. The elements chosen to act as references must be within the list of solved elements, for the first pass, the

initially passed references will be the only solved element(s) and will therefore be chosen. Then, elements to be included to the subdivision are selected based on maximizing connections to the chosen references. The number of elements to include in each subdivision is set by the function caller. The brute force algorithm is then applied to the subset of elements chosen to obtain alignment weights for that subdivision. One potential result of this process applied to an 800 element array is shown in Fig. 5.2.

As can be seen from the results, this process has some major imperfections that will need to be resolved in future research. There is a small subset of elements which are inaccurately calibrated, likely due to compounding errors from the chaining of subdivisions. These effects can be minimized through careful selection of the various parameters, but in a practical scenario this optimization cannot occur as there will be no ground truth to compare to. For the algorithm to be implemented in practice, the optimal setup configuration must be obtainable without trial and error.

The core difficulty that arises when designing such an algorithm lies in the computational complexity of determining if a given set of elements form a fully determined set of equations. While it may seem tempting to simply build the $\bar{\mathbf{A}}$ matrix and then compute the rank, that process is only marginally faster if faster at all than solving the system entirely. Similarly, many solutions that may seem obvious such as applying graph theory are equally computationally complex or will not actually give a metric of solution accuracy. Just because a particular element was included in the system of equations does not imply that it will be determined or increase the rank of the system. Rank alone is also not enough to determine if the system will solve to an accurate state.

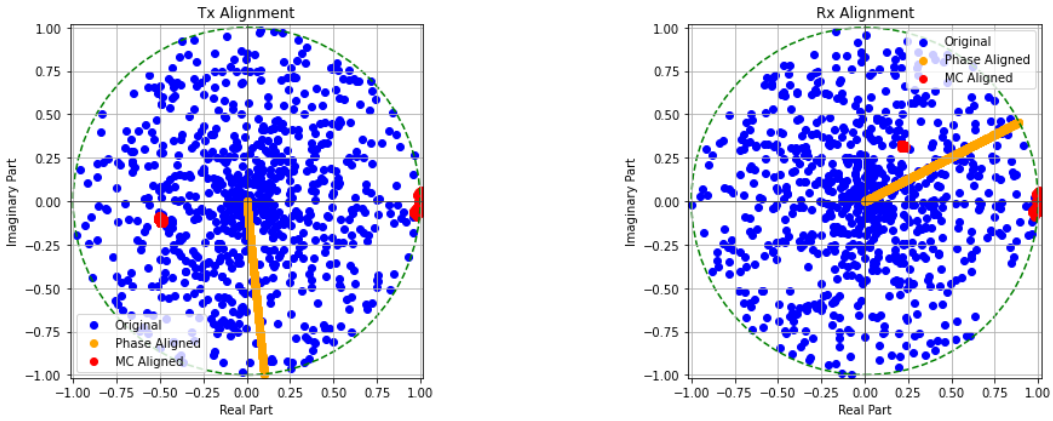

Figure 5.2: Results of connectivity based subdivision being applied to an 800 element array.

Chapter 6

Conclusion

Phased array radar systems rely on precise calibration to ensure optimal beamforming performance. Traditional calibration methods, such as those utilizing external reference antennas and near-field scanners, provide high accuracy but are often impractical for in-field calibration. Mutual coupling-based approaches have emerged as a promising alternative, leveraging the inherent interactions between radiating elements to determine alignment weights. However, existing methods often depend heavily on known array geometries, making them less adaptable to real-world constraints such as amplifier saturation, missing data, and imperfect coupling measurements.

This research proposes multiple geometry agnostic solutions that may provide higher accuracy results for certain applications. Ultimately, the proposed framework not only generalizes existing mutual coupling calibration methods but also enables their seamless integration within a unified computational model. This flexibility allows for improved performance and reliability in a wide range of practical radar applications, contributing to the advancement of phased array technology.

Appendix A

The Moore-Penrose Psuedo-Inverse

A.1 Definition

The Moore-Penrose pseudoinverse is a generalization of the matrix inverse that can be applied to non-square or singular matrices. For a given matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$, the pseudoinverse, denoted as $\mathbf{A}^{\dagger}$, satisfies the following four Penrose conditions:

1. $\mathbf{A}\mathbf{A}^{\dagger}\mathbf{A} = \mathbf{A}$,
2. $\mathbf{A}^{\dagger}\mathbf{A}\mathbf{A}^{\dagger} = \mathbf{A}^{\dagger}$,
3. $(\mathbf{A}\mathbf{A}^{\dagger})^{T} = \mathbf{A}\mathbf{A}^{\dagger}$,
4. $(\mathbf{A}^{\dagger}\mathbf{A})^{T} = \mathbf{A}^{\dagger}\mathbf{A}$.

A.2 Computation Methods

There are several numerical methods to compute the Moore-Penrose pseudoinverse, with the most commonly used being Singular Value Decomposition (SVD) and QR decomposition.

Singular Value Decomposition (SVD)

Given a matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$, the SVD is expressed as:

$$\mathbf{A} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T,$$

where:

- $\mathbf{U} \in \mathbb{R}^{m \times m}$ is an orthogonal matrix,
- $\mathbf{V} \in \mathbb{R}^{n \times n}$ is an orthogonal matrix,
- $\mathbf{\Sigma} \in \mathbb{R}^{m \times n}$ is a diagonal matrix with singular values σ_i on the diagonal.

The pseudoinverse is then computed as:

$$\mathbf{A}^\dagger = \mathbf{V}\mathbf{\Sigma}^\dagger\mathbf{U}^T,$$

where $\mathbf{\Sigma}^\dagger$ is obtained by taking the reciprocal of each non-zero singular value in $\mathbf{\Sigma}$ and transposing it. If $\sigma_i = 0$, then the reciprocal is set to zero to avoid division by zero.

A.2.1 QR Decomposition

Another approach to compute the pseudoinverse is using the QR decomposition, which is particularly effective for matrices that have full column rank. For a matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$, if $m \geq n$, the QR decomposition is expressed as:

$$\mathbf{A} = \mathbf{QR},$$

where:

- $\mathbf{Q} \in \mathbb{R}^{m \times n}$ is an orthogonal matrix,
- $\mathbf{R} \in \mathbb{R}^{n \times n}$ is an upper triangular matrix.

The pseudoinverse is then computed using:

$$\mathbf{A}^{\dagger} = \mathbf{R}^{-1}\mathbf{Q}^{T},$$

assuming $\mathbf{A}$ has full column rank and $\mathbf{R}$ is invertible.

A.3 Applications

The Moore-Penrose pseudoinverse is widely used in solving linear least-squares problems, where it provides a way to find the minimum norm solution for systems of linear equations that may not have a unique solution. It is also used in machine learning algorithms, control theory, and signal processing. The simplest way to implement this algorithm is with the np.linalg.pinv() function in python (or the pinv() function in Matlab). However, for applications which require many iterative loops, the numpy function may be inefficient.

A.4 Fast Computation of the Moore-Penrose Inverse

The algorithm used for the fast computation of Moore-Penrose inverse matrices, as proposed by Courrieu [37] can be used to increase the speed of computation for the algorithms described in this document. The downside is that this implementation is not stable for all matrices, however, in practice it does not appear to be unstable for mutual coupling datasets tested for this research. This method is employed throughout our paper to efficiently handle large, potentially rank-deficient matrices.

A.4.1 Algorithm Description

Let G be an $m \times n$ real matrix. The Moore-Penrose inverse G^+ is defined as the unique $n \times m$ matrix satisfying the following four properties:

$$GG^+G = G, \tag{A.1}$$

$$G^+GG^+ = G^+, \tag{A.2}$$

$$(GG^+)' = GG^+, \tag{A.3}$$

$$(G^+G)' = G^+G. \tag{A.4}$$

The standard approach to compute G^+ involves Singular Value Decomposition (SVD), which, while accurate, can be computationally expensive for large matrices. Courrieu's method leverages a full-rank Cholesky factorization to significantly reduce computation time.

A.4.2 Algorithm Foundations

Consider the symmetric positive semi-definite matrix $G'G$. If $G'G$ has rank $r \leq n$, there exists a unique upper triangular matrix S with exactly $(n-r)$ zero rows such that:

$$S'S = G'G. \tag{A.5}$$

By removing the zero rows from S, we obtain an $r \times n$ full-rank matrix L', satisfying:

$$G'G = S'S = LL'. \tag{A.6}$$

Using the general relation for the Moore-Penrose inverse of a matrix product [38]:

$$(AB)^+ = B'(A'ABB')^+A', \tag{A.7}$$

and setting $B = I$, we have:

$$A^+ = (A'A)^+A'. \tag{A.8}$$

Applying this to our matrix G, we derive:

$$G^+ = (G'G)^+G' = (LL')^+ G'. \tag{A.9}$$

Using the result for $(AA')^+$ when A is of full rank:

$$(AA')^+ = A(A'A)^{-1}(A'A)^{-1}A', \tag{A.10}$$

we obtain the final expression:

$$G^{+} = L(L'L)^{-1}(L'L)^{-1}L'G'. \tag{A.11}$$

A.4.3 Implementation Details

The algorithm consists of two main operations:

1. Full-rank Cholesky factorization of $G'G$ to obtain L.

2. Inversion of the symmetric positive definite matrix $L'L$.

On serial processors, these operations have computational complexities of $\mathcal{O}(n^3)$ and $\mathcal{O}(r^3)$, respectively. However, in parallel architectures, the time complexity can be significantly reduced, with Cholesky factorization potentially achieving $\mathcal{O}(n)$ and matrix inversion $\mathcal{O}(\log r)$ [39].

The MATLAB implementation of this algorithm is provided in Courrieu's original paper, demonstrating its efficiency and accuracy compared to traditional methods.

Appendix B

Limitation of Logarithmic Representation

The natural logarithm of a complex number $z = x + jy$ (where x and y are real numbers, and j is the imaginary unit) is defined using the polar form of the complex number. A complex number z can be expressed in polar form as:

$$z = re^{j\theta},$$

where:

- $r = |z| = \sqrt{x^2 + y^2}$ is the magnitude of z,
- $\theta = \arg(z)$ is the phase of z, representing the angle in the complex plane.

The natural logarithm of z is then given by:

$$\ln(z) = \ln(re^{j\theta}) = \ln(r) + j\theta.$$

Here, $\ln(r)$ is the natural logarithm of the magnitude, and θ is the phase of the complex number.

Phase Wrapping and the 90-Degree Constraint

The logarithm of a complex number is inherently multi-valued due to the periodic nature of the complex exponential function. Specifically, the phase θ is only defined up to integer multiples of 2π. That is, θ and $\theta + 2\pi k$ (for any integer k) represent the same complex number. This means that:

$$\ln(z) = \ln(r) + j(\theta + 2\pi k),$$

where k is an integer. The principal value of the logarithm is typically defined by restricting θ to the interval $(-\pi, \pi]$ (or equivalently, $-180°$ to $180°$).

When we transfer the calibration problem to the complex logarithmic domain, we are effectively working with the logarithms of the transmit (Tx) and receive (Rx) module coefficients. However, this transformation introduces a critical limitation: the phase differences between elements must be small enough to avoid phase wrapping ambiguities. Specifically, the phase variation of the Tx and Rx modules must be constrained to $\pm 90°$ to ensure that the system of equations can be solved uniquely.

Why the 90-Degree Condition Applies

The $\pm 90°$ constraint arises because phase wrapping ambiguities occur when the phase difference ($\Delta\theta$) between two elements exceeds $\pm 180°$. By constraining the phase variation of each Tx and Rx module to $\pm 90°$, we ensure that the sum or difference of phases (e.g., $\theta_1 + \theta_2$ or $\theta_1 - \theta_2$) never exceeds $\pm 180°$. This prevents phase wrapping entirely and ensures that the logarithmic representation remains unambiguous.

For example, consider two complex numbers $z_1 = e^{j\theta_1}$ and $z_2 = e^{j\theta_2}$. The phase difference between them is $\Delta\theta = \theta_2 - \theta_1$. If $\Delta\theta$ exceeds $\pm 180°$, it becomes difficult to determine whether the phase difference is $\Delta\theta$ or $\Delta\theta \pm 2\pi$. This ambiguity disrupts the linearization of the calibration equations, making it impossible to solve for the module coefficients accurately.

By constraining the phase variation to $\pm 90°$, we ensure that:

$$|\theta_1 - \theta_2| \leq 180°,$$

even in the worst-case scenario where θ_1 and θ_2 are at opposite extremes of their allowed ranges. This guarantees that phase wrapping cannot occur, and the logarithmic representation remains valid.

Implications for Calibration

The $\pm 90°$ constraint is a practical limitation imposed to ensure that the logarithmic representation remains unambiguous. To address this, a pre-alignment step is necessary to ensure that the phase variation of the Tx and Rx modules remains within the $\pm 90°$ range. This pre-alignment step involves applying a phase alignment algorithm to correct for phase wrapping effects before proceeding with the calibration process. By doing so, we ensure that the logarithmic representation remains valid and that the calibration algorithm can produce accurate results.

References

[1] C. Fulton *et al.*, "Mutual coupling-based calibration for the Horus digital phased array radar," in *2022 IEEE International Symposium on Phased Array Systems & Technology (PAST)*, Waltham, MA, USA, 2022, pp. 1–6. DOI: 10.1109/PAST49659.2022.9974962.

[2] E. Langley, C. Fulton, and M. Yeary, "Analysis of mutual coupling algorithms for a dipole array," in *2021 IEEE International Conference on Microwaves, Antennas, Communications and Electronic Systems (COMCAS)*, Tel Aviv, Israel, 2021, pp. 326–329. DOI: 10.1109/COMCAS52219.2021.9629046.

[3] C. Fulton, J. Salazar, D. Zrnic, D. Mirkovic, I. Ivic, and D. Doviak, "Polarimetric phased array calibration for large-scale multi-mission radar applications," in *2018 IEEE Radar Conference (RadarConf18)*, Oklahoma City, OK, USA, 2018, pp. 1272–1277. DOI: 10.1109/RADAR.2018.8378746.

[4] D. Bekers, R. van Dijk, and F. van Vliet, "Mutual-coupling based phased-array calibration: A robust and versatile approach," in *2013 IEEE International Symposium on Phased Array Systems and Technology*, IEEE, 2013, pp. 630–637.

[5] K. Hassett, "Phased array antenna calibration measurement techniques and methods," in *European Conference Antennas Propagation (EuCAP)*, 2016.

[6] A. E. Mitchell, "Coupling-based wideband digital phased array calibration techniques," *University of Oklahoma MS Thesis*, 2014.

[7] C. Fulton, P. Kenworthy, J. Lujan, M. Herndon, S. Garner, D. Thompson, and M. Yeary, "Mutual coupling-based calibration for the Horus digital phased array radar," in *2022 IEEE International Symposium on Phased Array Systems & Technology (PAST)*, IEEE, 2022, pp. 1–6.

[8] H. M. Aumann, A. J. Fenn, and F. G. Willwerth, "Phased array antenna calibration and pattern prediction using mutual coupling measurements," *IEEE Transactions on Antennas and Propagation*, vol. 37, no. 7, pp. 844–850, 1989.

[9] O. Inac, S. Y. Kim, D. Shin, C. Y. Kim, and G. M. Rebeiz, "Built-in self test systems for silicon-based phased arrays," in *2012 IEEE/MTT-S International Microwave Symposium Digest*, Montreal, QC, Canada, 2012, pp. 1–3. DOI: 10.1109/MWSYM.2012.6259455.

[10] J. Lujan, C. J. Fulton, M. Yeary, E. A. Langley, S. McCormick, and A. Hedden, "Phased array radar initial alignment algorithm using mutual coupling: an iterative approach," in *Radar Sensor Technology XXIV*, K. I. Ranney and A. M. Raynal, Eds., International Society for Optics and Photonics, vol. 11408, SPIE, 2020, p. 1 140 811. DOI: 10.1117/12.2562963. [Online]. Available: https://doi.org/10.1117/12.2562963.

[11] I. Şeker, "Calibration methods for phased array radars," in *Proceedings of SPIE*, vol. 8714, May 2013, 87140W-1–87140W-15.

[12] R. M. Lebrón, P.-S. Tsai, J. M. Emmett, C. Fulton, and J. L. Salazar-Cerreno, "Validation and testing of initial and in-situ mutual coupling-based calibration of a dual-polarized active phased array antenna," *IEEE Access*, vol. 8, pp. 78 315–78 329, 2020. DOI: 10.1109/ACCESS.2020.2983523.

[13] H. Singh and R. M. Jha, "Mutual coupling effects in phased arrays," in *Active Radar Cross Section Reduction: Theory and Applications*. Cambridge University Press, 2015, pp. 177–215.

[14] Z. Dunn, M. Yeary, C. Fulton, and N. Goodman, "Wideband digital predistortion of solid-state radar amplifiers," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 52, no. 5, pp. 2452–2466, Oct. 2016. DOI: 10.1109/TAES.2016.150142.

[15] C. A. Balanis, *Antenna Theory Analysis and Design*. John Wiley & Sons, Inc., 2005.

[16] *HFSS getting starting guides/radar cross section (RCS) model*, HFSS, 2015.

[17] J. L. Salazar, N. Aboserwal, J. D. Díaz, J. A. Ortiz, and C. Fulton, "Edge diffractions impact on the cross polarization performance of active phased array antennas," in *2016 IEEE International Symposium on Phased Array Systems and Technology (PAST)*, 2016, pp. 1–5. DOI: 10.1109/ARRAY.2016.7832571.

[18] G. F. Lewis and E. Boe, *Self-phase up of array antennas with non-uniform element mutual coupling and arbitrary lattice orientation*, US Patent 5,657,023, Aug. 1997.

[19] A. Agrawal and A. Jablon, "A calibration technique for active phased array antennas," in *IEEE International Symposium on Phased Array Systems and Technology, 2003.*, 2003, pp. 223–228. DOI: 10.1109/PAST.2003.1256985.

[20] R. D. Palmer *et al.*, "Horus—a fully digital polarimetric phased array radar for next-generation weather observations," *IEEE Transactions on Radar Systems*, vol. 1, pp. 96–117, 2023. DOI: 10.1109/TRS.2023.3280033.

[21] M. Yeary, R. Palmer, C. Fulton, J. Salazar, and H. Sigmarsson, "Update on an S-band all-digital mobile phased array radar," in *2021 IEEE Radar Conference (RadarConf21)*, 2021, pp. 1–4.

[22] Y. Neidman, R. Shavit, and A. Bronshtein, "Diagnostic of phased arrays with fault elements using the mutual coupling method," in *2008 IEEE International Conference on Microwaves, Communications, Antennas and Electronic Systems*, 2008, pp. 1–6. DOI: 10.1109/COMCAS.2008.4562802.

[23] T. Takahashi, Y. Konishi, S. Makino, H. Ohmine, and H. Nakaguro, "Fast measurement technique for phased array calibration," *IEEE Transactions on Antennas and Propagation*, vol. 56, no. 7, pp. 1888–1899, 2008. DOI: 10.1109/TAP.2008.924682.

[24] R. Long, J. Ouyang, F. Yang, W. Han, and L. Zhou, "Multi-element phased array calibration method by solving linear equations," *IEEE Transactions on Antennas and Propagation*, vol. 65, no. 6, pp. 2931–2939, 2017. DOI: 10.1109/TAP.2017.2694767.

[25] A. C. Newell, B. Schluper, and R. J. Davis, "Holographic projection to an arbitrary plane from spherical near-field measurements," in *Sympo-*

sium Digest of the Antenna Measurement Techniques Association, 2001, pp. 92–97.

[26] D. Schvartzman, J. D. D. Díaz, D. Zrnić, M. Herndon, M. B. Yeary, and R. D. Palmer, "Holographic back-projection method for calibration of fully digital polarimetric phased array radar," *IEEE Transactions on Radar Systems*, vol. 1, pp. 295–307, 2023. DOI: 10.1109/TRS.2023.3286849.

[27] D. Schvartzman *et al.*, "A polarimetric antenna-calibration method for the Horus radar based on E-field back projection," in *2022 IEEE International Symposium on Phased Array Systems & Technology (PAST)*, Waltham, MA, USA, 2022, pp. 01–07. DOI: 10.1109/PAST49659.2022.9975030.

[28] A. Y. Umeyama, J. L. Salazar-Cerreno, and C. J. Fulton, "UAV-based far-field antenna pattern measurement method for polarimetric weather radars: Simulation and error analysis," *IEEE Access*, vol. 8, pp. 191 124–191 137, 2020. DOI: 10.1109/ACCESS.2020.3027790.

[29] J. L. Salazar-Cerreno, S. S. Jehangir, N. Aboserwal, A. Segales, and Z. Qamar, "An UAV-based polarimetric antenna measurements for radar and communication systems from 3 GHz to 32 GHz," in *2021 IEEE Conference on Antenna Measurements & Applications (CAMA)*, Antibes Juan-les-Pins, France, 2021, pp. 55–60. DOI: 10.1109/CAMA49227.2021.9703660.

[30] A. Y. Umeyama, J. L. Salazar-Cerreño, B. M. Wolf, and C. J. Fulton, "Recent development in UAV-based antenna pattern characterization for weather radars," in *2019 IEEE Conference on Antenna Measurements & Applications (CAMA)*, Kuta, Bali, Indonesia, 2019, pp. 199–202. DOI: 10.1109/CAMA47423.2019.8959698.

[31] A. Y. Umeyama, J. L. Salazar-Cerreno, and C. J. Fulton, "UAV-based antenna measurements for polarimetric weather radars: Probe analysis," *IEEE Access*, vol. 8, pp. 191 862–191 874, 2020. DOI: 10.1109/ACCESS.2020.3027779.

[32] A. Nafe, K. Kibaroglu, M. Sayginer, and G. M. Rebeiz, "An in-situ self-test and self-calibration technique utilizing antenna mutual coupling for 5g multi-beam TRX phased arrays," in *2019 IEEE MTT-S*

International Microwave Symposium (IMS), 2019, pp. 1229–1232. DOI: 10.1109/MWSYM.2019.8701072.

[33] C. Shipley and D. Woods, "Mutual coupling-based calibration of phased array antennas," in *Proceedings 2000 IEEE International Conference on Phased Array Systems and Technology (Cat. No.00TH8510)*, 2000, pp. 529–532. DOI: 10.1109/PAST.2000.859012.

[34] C. Fulton and W. Chappell, "Calibration techniques for digital phased arrays," in *2009 IEEE International Conference on Microwaves, Communications, Antennas and Electronics Systems*, 2009, pp. 1–10. DOI: 10.1109/COMCAS.2009.5385979.

[35] D. R. Morgan, Z. Ma, J. Kim, M. G. Zierdt, and J. Pastalan, "A generalized memory polynomial model for digital predistortion of RF power amplifiers," *IEEE Transactions on Signal Processing*, vol. 54, no. 10, pp. 3852–3860, Oct. 2006. DOI: 10.1109/TSP.2006.879264.

[36] M. Herndon and M. Yeary, "Real-time FPGA-based digital predistortion for improved amplifier performance in next generation phased arrays," in *2022 IEEE Radar Conference (RadarConf22)*, New York City, NY, USA, 2022, pp. 1–6. DOI: 10.1109/RadarConf2248738.2022.9764300.

[37] P. Courrieu, "Fast computation of Moore-Penrose inverse matrices," *Neural Information Processing - Letters and Reviews*, vol. 8, no. 2, pp. 25–29, 2005.

[38] A. Ben-Israel and T. N. E. Greville, *Generalized Inverses: Theory and Applications*, 2nd. Springer, 2003.

[39] P. E. Gill, W. Murray, and M. H. Wright, *Numerical Linear Algebra and Optimization*. Addison-Wesley, 1991.

[40] C. A. Balanis, *Advanced Engineering Electromagnetics*. John Wiley & sons, Inc., 2012.