

UNIVERSITY OF OKLAHOMA
GRADUATE COLLEGE

COUPLING-BASED WIDEBAND DIGITAL PHASED ARRAY CALIBRATION
TECHNIQUES

A THESIS

SUBMITTED TO THE GRADUATE FACULTY

in partial fulfillment of the requirements for the

Degree of

MASTER OF SCIENCE

By

ADAM E. MITCHELL
Norman, Oklahoma
2014

OU
THESIS
MIT
COP. 2

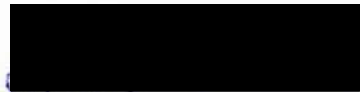
COUPLING-BASED WIDEBAND DIGITAL PHASED ARRAY CALIBRATION
TECHNIQUES

A THESIS APPROVED FOR THE
SCHOOL OF ELECTRICAL AND COMPUTER ENGINEERING

BY



Dr. Caleb Fulton, Chair



Dr. Hjalti Sigmarsson



Dr. Jessica Ruyle

Acknowledgements

I would like to thank Dr. Heidi Sigurdson and Dr. Jessica Rogie for their guidance on this project and all their instruction as I have begun my career in the field of HR. This work would not have been possible without the direction of my advisor, Dr. Cathie Palmer, and the financial support and expertise of Northington Crossman.

Acknowledgements

I would like to thank Dr. Hjalti Sigmarsson and Dr. Jessica Ruyle for their guidance on this thesis and all their instruction as I have begun my career in the field of RF. This work would not have been possible without the direction of my advisor, Dr. Caleb Fulton, or the financial support and expertise of Northrup Grumman.

1.1: The Nonlinear Calibrations	1
1.2: Mutual Coupling-based Methodologies	2
1.3: Vector Array Geometry	3
1.4: In-band versus In-Site Calibrations	3
1.5: The TR Matrix	3
Chapter 2: Software Framework	4
2.1: Software Components	4
2.1.1: RadModel	5
2.1.2: AntEmul	7
2.1.3: ChannelEmul	8
2.1.4: LinAnalysis	10
2.1.5: Interf	11
2.1.6: Engine	12
2.2: Process Flow	13
Chapter 3: In-band Calibration	14
3.1: Measurement Models	14
3.2: Software Implementation	18
3.1.1: Coupling Set Analysis	19

Table of Contents

Acknowledgements	iv
List of Figures.....	vii
Abstract.....	xii
Chapter 1: Introduction.....	1
1.1: The Need for Calibration.....	1
1.2: Mutual Coupling-based Methodologies	2
1.3: Test Array Geometry	3
1.4: Initial versus In-Situ Calibration	4
1.5: The TR Matrix.....	5
Chapter 2: Software Framework	7
2.1: Software Components	7
2.1.1: RadModel	9
2.1.2: AddError.....	9
2.1.3: ChannelError	10
2.1.4: LinAnalysis	10
2.1.5: Interf	11
2.1.6: Engine.....	12
2.2: Process Flow.....	12
Chapter 3: Initial Calibration.....	14
3.1: Mathematical Basis	14
3.2: Software Implementation	18
3.3: Coupling Set Analysis	19

3.4: Results	24
Chapter 4: In-Situ Calibration	35
4.1: Mathematical Basis	35
4.2: Combining Multiple Iterations	36
4.3: Software Implementation	37
4.4: Results	38
4.5: Tracking S-Parameter Errors.....	43
Chapter 5: Calibration Method Comparison	46
5.1: Test Setup.....	46
5.2: Results	46
Chapter 6: Conclusion	52
References	54
Appendix A: Populating the TR Matrix	56

List of Figures

Fig. 1: Example of linear array indexing with the active five by five subset used in testing highlighted in green	4
Fig. 2: Block diagram illustrating the software framework	8
Fig. 3: Block diagram illustrating the sequential order of function calls.	13
Fig. 4: Overlapping plots of the coupling values (magnitude in dB) within each set at every frequency	21
Fig. 5: The standard deviation of the coupling values within each set at every frequency.	22
Fig. 6: The standard deviation of the ratios of coupling values within each set across frequency.	22
Fig. 7: Overlapping plots of the ratios of coupling values within each set at every frequency	23
Fig. 8: Frequency response of element 23's transmit chain when ideal s-parameters are used. A .005 offset has been added for the sake of visibility.	25
Fig. 9: (Left) The altered center reference element response normalized to itself, in red, and normalized to the nominal reference response, in blue, (Right) All other channel responses relative to the altered reference, in red, and relative to the nominal reference channel, in blue.....	26
Fig. 10: Frequency response of element 43's transmit chain when simulated s-parameters are used. Random complex error has been added to all elements except the center element, which is used as the reference.....	27

Fig. 11: Frequency response of element 21's transmit chain when simulated s-parameters are used. Random complex error has been added to all elements except the center element, which is used as the reference. 28

Fig. 12: Frequency response of element 57's transmit chain when simulated s-parameters are used. Random complex error has been added to all elements except the center element, which is used as the reference. 29

Fig. 13: Results at 3.17 GHz when a single element is manually altered, (left) Estimated element gain magnitudes, (center) Actual element gain magnitudes, (right) Magnitude of the difference between estimated and actual gain magnitudes. 30

Fig. 14: Results at 3.17 GHz when a single element is manually altered, (left) Estimated element gain phases at 3.17 GHz, (center) Actual element gain phases, (right) Magnitude of the difference between estimated and actual gain phases. 30

Fig. 15: Results at 3.09 GHz when a single element is randomly altered, (left) Estimated element gain magnitudes, (center) Actual element gain magnitudes, (right) Magnitude of the difference between estimated and actual gain magnitudes 30

Fig. 16: Results at 3.3 GHz when elements are randomly altered, (left) Estimated element gain magnitudes, (center) Actual element gain magnitudes, (right) Magnitude of the difference between estimated and actual gain magnitudes 31

Fig. 17: Results at 3.3 GHz when all but the center element are randomly altered, (left) Estimated element gain phases, (center) Actual element gain phases, (right) Magnitude of the difference between estimated and actual gain phases 31

Fig. 18: A comparison of the original absolute magnitude and phase differences from [12] with those calculated using the reverse engineered methodology 32

Fig. 19: Differences between estimated and actual amplitude and phase at 3.3 GHz when the subset is placed in the corner	33
Fig. 20: Differences between estimated and actual amplitude and phase at 3.55 GHz when the subset is placed in the corner	34
Fig. 21: Differences between estimated and actual amplitude and phase at 3.3 GHz when the subset is placed in the center.....	34
Fig. 22: Differences between estimated and actual amplitude and phase at 3.55 GHz when the subset is placed in the center.....	34
Fig. 23: The five transmit/receive configurations used to calculate the complete array error. The red represents transmitting elements, the blue represents receiving elements, and gray elements are unused.....	36
Fig. 24: (left) Estimated frequency response of element 60's receive chain in blue and the actual response in red, (right) Overlapping plots of the difference between the estimated and actual responses for every element in the receive group.....	39
Fig. 25: (left) Estimated frequency response of element 34's transmit chain in blue and the actual response in red, (right) Overlapping plots of the difference between the estimated and actual responses for every element in the transmit group	39
Fig. 26: (left) Estimated frequency response of element 34's transmit chain in blue and the actual response with a .01 offset in red, (right) Overlapping plots of the difference between the estimated and actual responses for every element in the transmit group ...	40
Fig. 27: (left) Estimated frequency response of element 60's receive chain in blue and the actual response in red, (right) Estimated frequency response of element 49's receive chain in blue and the actual response in red	41

Fig. 28: (left) Estimated frequency response of element 34's transmit chain in blue and the actual response in red, (right) Estimated frequency response of element 22's transmit chain in blue and the actual response in red.....	42
Fig. 29: (left) Overlapping plots of the difference between the estimated and actual responses for every element in the receive group, (right) Overlapping plots of the difference between the estimated and actual responses for every element in the transmit group.....	42
Fig. 30: The normalized singular value sum, α , plotted versus the error scale factor. Plot shown for 3.3	44
Fig. 31: The normalized singular value sum, α , plotted versus the RMS element error. Plot shown for 3.3 GHz.....	45
Fig. 32: RMS Element Error from the initial calibration routine as a function of the error scales.....	47
Fig. 33: RMS Element Error from the in-situ calibration routine as a function of the error scales.....	47
Fig. 34: In-situ calibration errors when the golden standard has an error scale of .1 applied	49
Fig. 35: In-situ calibration errors when the golden standard has an error scale of .2 applied	49
Fig. 36: In-situ calibration errors when the golden standard has an error scale of .3 applied	50
Fig. 37: Average of in-situ calibration errors from ten runs when the golden standard has an error scale of .4 applied	50

Fig. 38: Element 59's transmit channel frequency response when the golden standard has an error scale of .4, the measurement's s-parameters have an error scale of .2, and there is no gain error 51

Fig. 39: Example of a two by two array with red representing transmit, blue representing receive, and gray representing unused elements..... 56

Fig. 40: Illustration of how the full coupling matrix is reduced..... 56

Abstract

Recently, interest in phased array radars for diverse applications has been rising because of their versatility and recent efforts to reduce their cost. All phased array radars require some form of calibration to periodically remove the amplitude and phase errors which arise during normal operation, regardless of the cost or design of the array. This thesis offers an analysis of two of the latest mutual coupling calibration techniques, one in-situ calibration method and one initial calibration method. The initial calibration mathematics that were reverse engineered from the original publication are presented. The in-situ calibration approach is expanded to allow calibration of an entire array face. A flexible software framework was designed to test both calibration techniques on practical simulation data and to offer a direct comparison of the results, bringing insight into the relative merits of each approach.

Chapter 1: Introduction


1.1: The Need for Calibration

As advances continue to drive down the costs associated with phased array radar systems, the need for versatile and efficient calibration methods has increased. All active phased arrays require some form of calibration in order to ensure the amplitudes and phases of the elements are properly aligned. Errors in amplitude and phase can occur for any number of reasons, such as temperature changes, errors in the active circuitry, or mechanical deformation. These effects are more pronounced when the arrays are smaller and components of lesser quality are used, so the performance required of modern phased arrays would be impossible on these systems without some means of identifying errors. Over the years, countless calibration techniques have been developed. In [1] and [2], the elements are characterized by optimizing a log-likelihood function generated using the fields' known covariances. Clutter has also been used in conjunction with synthetic aperture radar (SAR) to generate measurements and calibrate element gains and phases [3]. Ground backscatter of MIMO signals can also be used to simultaneously estimate the direction-dependent transmit and receive calibration factors [4]. Direction-dependent calibration, like that in [4], accounts for the radiating properties of the antenna elements when producing calibration factors. It is worth noting that all of the above approaches rely on the systems' ability to receive a signal from a far off source or scatterer. This work will focus on the utilization of mutual coupling to calibrate both transmitters and receiver, and it will concentrate on correcting at the channel level. Here, the term channel refers to the active RF circuitry used to excite the array. This should not be confused with the common communications theory

terminology where channel refers to the transmission medium. The addition of direction-dependence is left for future work. Mutual coupling is generally defined as the ratio of the received signal to the transmitted signal [5]. Relying on interelement coupling simplifies the process of taking measurements with multiple transmitters and receivers at a time. This can have an averaging effect, improving the accuracy of the measurements. Mutual coupling calibration also requires less hardware than using dedicated calibration networks, like those in [6] and [7]. Calibration networks are immune to near field effects, and it is possible to use wave coding techniques to further improve performance. However, they often require additional complicated timing or, as is the case in [6], matching circuitry. Additional care must also be taken to ensure the network remains properly calibrated. Mutual coupling eliminates the need for cumbersome external measurement equipment like that used in [8] and [9]. Both [8] and [9] attempt to use external equipment to measure the far field radiation; [8] utilizes near field measurements as well. Near field and far field measurements have proven invaluable when characterizing an array in rigorous laboratory conditions, but they are extremely sensitive to the reflections, leakages, and interference commonly encountered when an array must be calibrated in the field. Mutual coupling based calibration methods will be discussed further in Section 1.2.

1.2: Mutual Coupling-based Methodologies

Aumann was one of the first to investigate mutual coupling based calibration [5]. He demonstrated a way to correct for module amplitude and phase errors and predict the array radiation pattern based on the assumption that elements were uniformly spaced and had identical, symmetric radiation patterns. The methodology



also required the array be large (thousands of elements) to ensure the mutual coupling between adjacent elements did not depend on the pair's position within the array. Antenna pairs' identical couplings meant large groups of measurements could be compared, and the resulting differences had to be a result of module errors. A similar approach was used in [10], but it was shown that proper alignment could still be achieved in smaller arrays (hundreds of elements) with more irregular spacing by tying together elements within different measurement groups. In [11], Neidman illustrates mutual coupling methods' sensitivity to faulty elements and how errors will propagate throughout the array during calibration. Neidman also proposed a way to identify and bypass faulty elements based on the work in [5] and [10]. This is a small sample of the work done on mutual coupling calibration over the years, but this work will focus on the methods demonstrated in [12] and [13] because they represent state-of-the-art techniques. [5], [10], [11], and other mutual coupling calibration schemes all rely on transmitting with a single element at a time. [12] and [13] are unique in that they calibrate multiple transmit elements at once, making them the most viable starting points. This paper generalizes these approaches based on recent work and, in the process, improves them.

1.3: Test Array Geometry

To ensure comparability between calibration tests, a single array configuration is used throughout. A nine element by nine element planar array was simulated using the finite element method solver HFSS®, and its s-parameters were stored in the radModel object for testing. To help mitigate edge effects, a five by five subset of the array was treated as the active region for calibration. When referring to elements, they will be

identified by their linear index. Figure 1 illustrates this indexing pattern as well as highlights the active five by five subset.

1	10	19	28	37	46	55	64	73
2	11	20	29	38	47	56	65	74
3	12	21	30	39	48	57	66	75
4	13	22	31	40	49	58	67	76
5	14	23	32	41	50	59	68	77
6	15	24	33	42	51	60	69	78
7	16	25	34	43	52	61	70	79
8	17	26	35	44	53	62	71	80
9	18	27	36	45	54	63	72	81

Fig. 1: Example of linear array indexing with the active five by five subset used in testing highlighted in green

1.4: Initial versus In-Situ Calibration

In this work, the terms initial calibration and in-situ calibration will be used extensively. Initial calibration will refer to the process of aligning the elements' amplitudes and phases without prior knowledge of their values, and it is typically done before fielding the array. The method in [12] has been chosen for initial calibration testing here, and it will be described in detail in Chapter 3. In-situ calibration refers to the process of correcting amplitudes and phases to bring them in line with their ideal values after the system has been operational for a period of time. Ideally, the corrections are made with minimal user interaction while the array is still in the field. The methods

outlined in [13] are used for in-situ calibration and are described in Chapter 4. While calibration typically includes the correction of errors, this work will focus on the estimation of the amplitude and phase errors because correction is trivial once error values are properly identified. The correction methodology usually depends heavily on the hardware's means of channel correction as well.

1.5: The \mathbf{TR} Matrix

The \mathbf{TR} matrix was originally described in [13], and it is a critical component of both the initial and the in-situ calibration routines. \mathbf{TR} is used to represent the combined effects of coupling, gain, and leakage on the signals travelling between the transmitting element group and the receiving element group. It is defined as

$$\mathbf{TR} = \mathbf{G}^R (\mathbf{I} + \mathbf{L}^R) \mathbf{C} (\mathbf{I} + \mathbf{L}^T) \mathbf{G}^T. \quad (1)$$

\mathbf{G}^T and \mathbf{G}^R are diagonal matrices containing the gains of the transmit and receive channels, respectively, but with the columns and rows corresponding to unused elements removed. \mathbf{I} is the identity matrix, and the \mathbf{L} terms represent the leakage between transmit and receive channels. For the digital arrays tested here, the leakage can be ignored because each channel has independent circuitry. Finally, the \mathbf{C} matrix is the coupling matrix in which the columns corresponding to unused transmit elements and the rows corresponding to unused receive elements have been removed. This yields the matrix \mathbf{TR} with the same unused rows and columns removed. It is assumed that all components are operating in their linear regimes. [13] maintains linearity by creating two separate sets of \mathbf{TR} measurements, one for the high power configurations and one for the low power configurations. Receiver errors are calculated using the low power version, and transmitter errors are calculated using the low power version. \mathbf{TR} can be

used to represent interelement couplings for a wide variety of transmit and receive group configurations, and it will play an important role in the calibration methods examined here. Appendix A further details the creation of a TR matrix.

Chapter 2: Software Framework

In lieu of being confined to a hardware platform for calibration tests, a flexible software framework has been developed to more fully explore the efficacy of different algorithms. This also enables one to test countless array configurations and real-world scenarios by simply changing the model parameters. For this work, everything has been implemented through a series of functions programmed using the MATLAB® software suite. Figure 2 illustrates the overall layout of the software platform. The rest of this chapter will go into greater detail about each of the blocks in Figure 2.

2.1: Software Components

Before diving into the larger workings of the framework, each component and its respective abilities must be understood. Figure 2 outlines some of these basic abilities within the light green boxes symbolizing separate MATLAB® files. Each of these blocks will now be explained heuristically, and the mathematical details will be fleshed out in later sections.

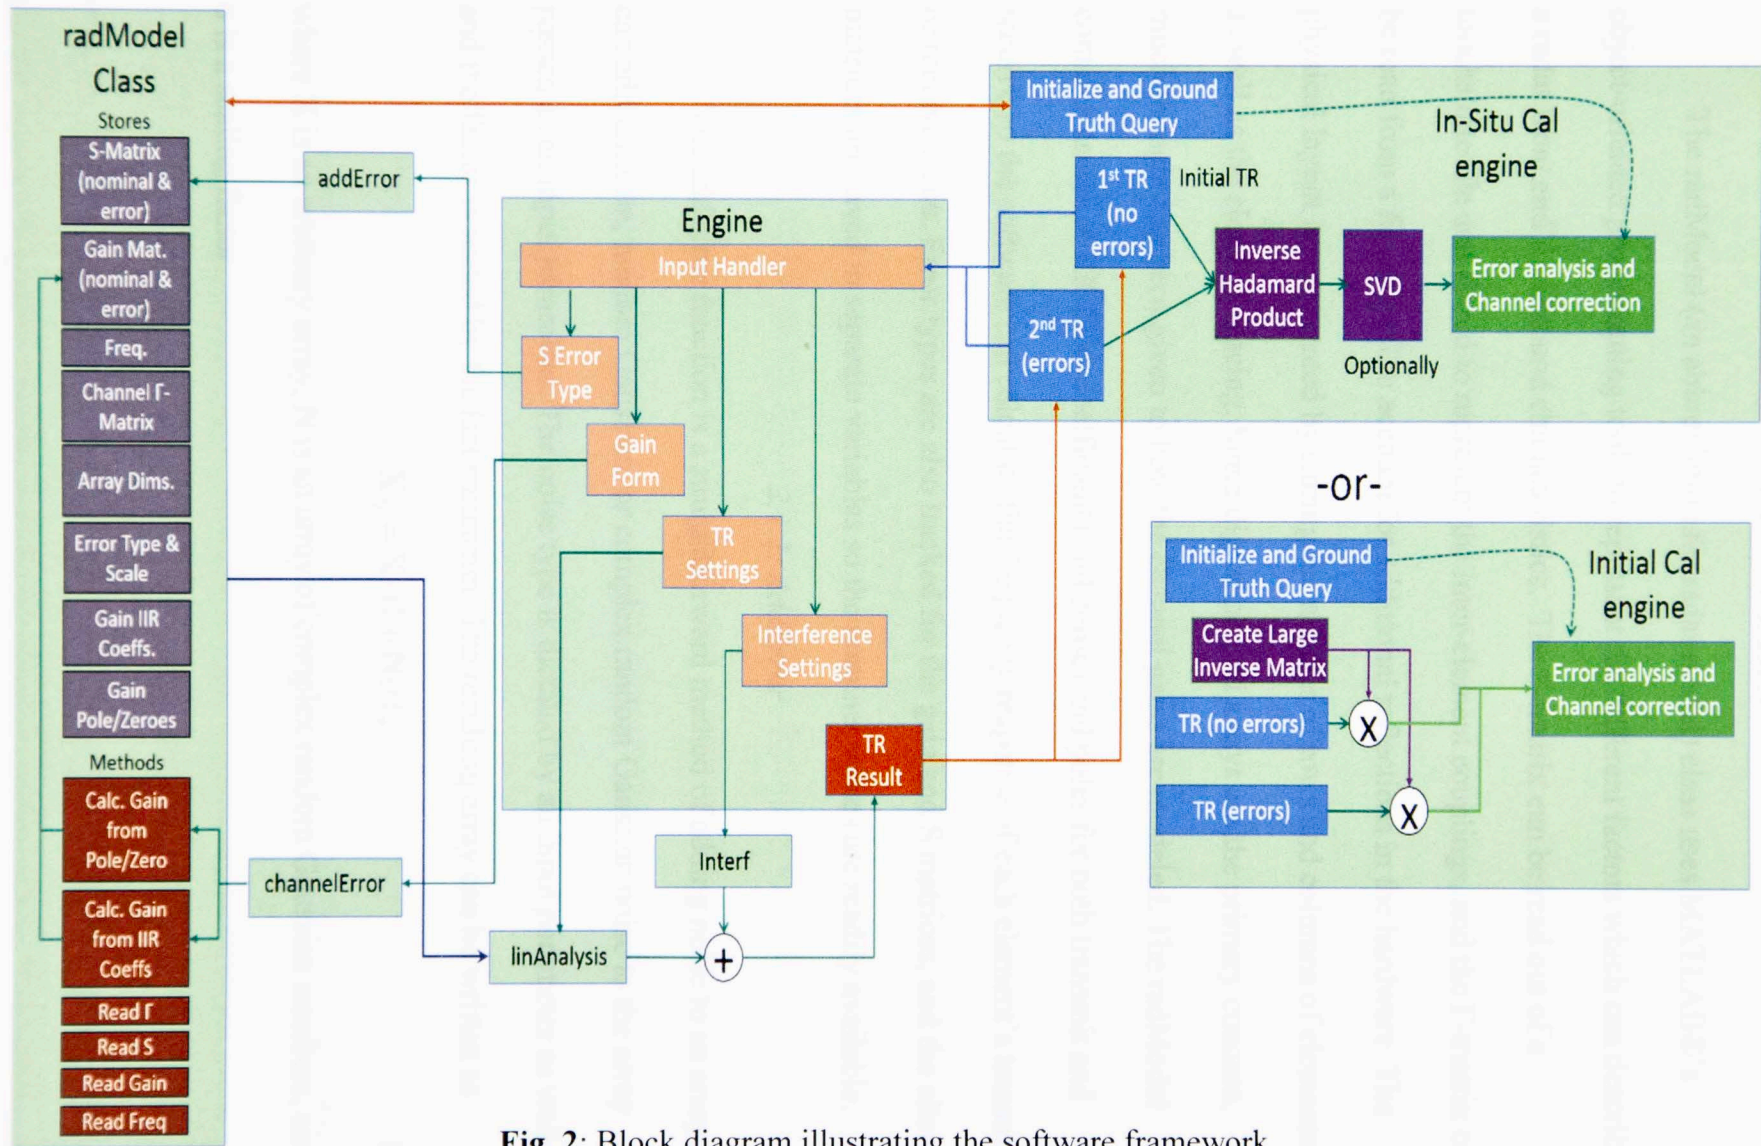


Fig. 2: Block diagram illustrating the software framework

2.1.1: RadModel

The radModel (an abbreviation for radar model) class uses MATLAB®'s object-oriented programming tools to represent the different factors which can describe a radar front end's wideband characteristics. The S-matrix can be read out of a touchstone file and stored to represent the inter-element couplings, and the Γ -matrix can be read from a file as well to account for additional reflections in the hardware. The physical layout is represented by storing the number of rows and columns of elements as well as the element spacing. Since digital phased arrays are the primary concern, much thought has been given to how the channel gains are modeled. The radModel contains arrays of IIR filter coefficients and zeroes and poles for both transmit and receive so the software can calculate the frequency response of each element's transmit or receive chain. Error types are also tracked for the gain and S matrices, and the altered matrices are saved in separate variables so the nominal values are readily available.

2.1.2: AddError

The addError function is a straightforward method of adding noise to an array. It can add uniform, random Gaussian, or complex random Gaussian noise to the array passed as an input parameter. The noise type is dictated by an input parameter as well, and the factor is scaled by the last parameter. The resulting array can be written as

$$\mathbf{X}_N = \mathbf{X} \cdot (1 + \mathbf{N}\sigma), \quad (2)$$

where \mathbf{X} is an arbitrary array, \mathbf{N} is an array of complex random Gaussian numbers, and σ is a scaling factor.

2.1.3: ChannelError

ChannelError alters the IIR coefficients and the zeroes and poles of the channels specified within the radModel. These adjusted values are then stored in separate radModel variables. The magnitude of these alterations depends upon the standard deviations stored in the radModel. The altered versions of the IIR coefficients, zeroes, and poles can be written as

$$\mathbf{A}_E = \mathbf{A} + \frac{\sigma_A}{\sqrt{2}} \mathbf{R}_A, \quad (3)$$

$$\mathbf{B}_E = \mathbf{B} + \frac{\sigma_B}{\sqrt{2}} \mathbf{R}_B, \quad (4)$$

$$\mathbf{Z}_E = \mathbf{Z} + \frac{\sigma_Z \mathbf{d}_Z}{\sqrt{2}} \mathbf{R}_Z, \quad (5)$$

$$\mathbf{P}_E = \mathbf{P} + \frac{\sigma_P \mathbf{d}_P}{\sqrt{2}} \mathbf{R}_P, \quad (6)$$

where \mathbf{A} , \mathbf{B} , \mathbf{Z} , and \mathbf{P} are matrices containing the IIR denominator, the IIR numerator, the zeroes, and the poles respectively. The \mathbf{R} term in each equation is a matrix of complex Gaussian random numbers, the σ terms represent the standard deviation of their respective parameters. The \mathbf{d} terms in equations five and six correspond to the poles' and zeroes' distance from the unit circle. Two versions of this function were used. The latest version was designed to more accurately reflect the behavior of errors in actual hardware. It will be indicated if the newer version of channelError was used.

2.1.4: LinAnalysis

LinAnalysis performs the linear algebra necessary to generate the TR matrix used for analysis. It will read any necessary values out of the radModel, and calculate the TR matrix for specific transmit and receive element configurations. Two additional

input parameters specify whether the gain and S-parameters are assumed to be nominal. The mathematics behind this function are briefly explained in Section 1.5, and a more rigorous explanation can be found in [13].

2.1.5: *Interf*

The function *interf* is designed to generate a matrix representing the noise generated at the receive elements by an external interference source. It is important to separate this type of noise from that generated in *addError* because the noise present at the face of adjacent receivers can be approximated as the same amplitude signal but phase shifted, meaning the noise is no longer uncorrelated. For each interferer, the noise signal at each element in the receive group is phase shifted based on the location of the interferer relative to the upper left-most receiver and the dimensions of the array. If the array is positioned on the *xz*-plane, the equation to calculate the noise at a receiver is

$$A \exp(k(x \sin(\theta) \cos(\varphi) + z \cos(\theta))), \quad (7)$$

where A is the amplitude of the noise source (modelled as a complex Gaussian random variable), k is the wave number, x and z are the Cartesian coordinates of the element's center relative to the center of the upper left receiver, and θ and φ are the spherical angles of the interference source. In the case of multiple interferers, the noise at each element is treated as the superposition of the noise from each interferer. The interferers are assumed to be statistically independent and are spectrally white, meaning that interference in each frequency bin is independent from that of other frequency bins. Furthermore, a given interferer's signal as it appears at the receivers during a coupling measurement from any given transmitter is assumed to be independent of the signal that

appears at the receivers during coupling measurements from a different transmitter. Under the assumption of spectrally white interference, this is valid both for coded transmitter schemes as well as time-separated transmitter measurements. In principle, the code could be modified to support deterministic transmit waveforms, but this would require more input from the user, including transmission timelines, coding schemes, etc. This is left for future work.

2.1.6: Engine

As Figure 2 shows, the engine function is how the user interacts with the various components of the system. Rather than the user call every function directly, he or she calls the engine with a variable input declaring how each function should run, and the engine ensures everything runs in the proper sequence with the correct settings. The engine function is capable of handling the following: calling `addError` to calculate the altered version of the S-matrix, calling `channelError` to change the gain model and calling the `radModel` functions to calculate the new frequency response, calling the `radModel` functions to set Γ from a file or set it to zero, calling `linAnalysis` to generate the TR matrix with nominal or altered versions of the matrices specified, and calling `interf` and adding the resulting matrix to the TR matrix generated by `linAnalysis`.

2.2: Process Flow

Figure 3 demonstrates the sequence of function calls which takes place before the calibration algorithms are used. The calibration algorithms would come after the end of this sequence and will be described in more detail in Chapters 3 and 4. Figure 3 helps elucidate the nested manor in which the functions operate.

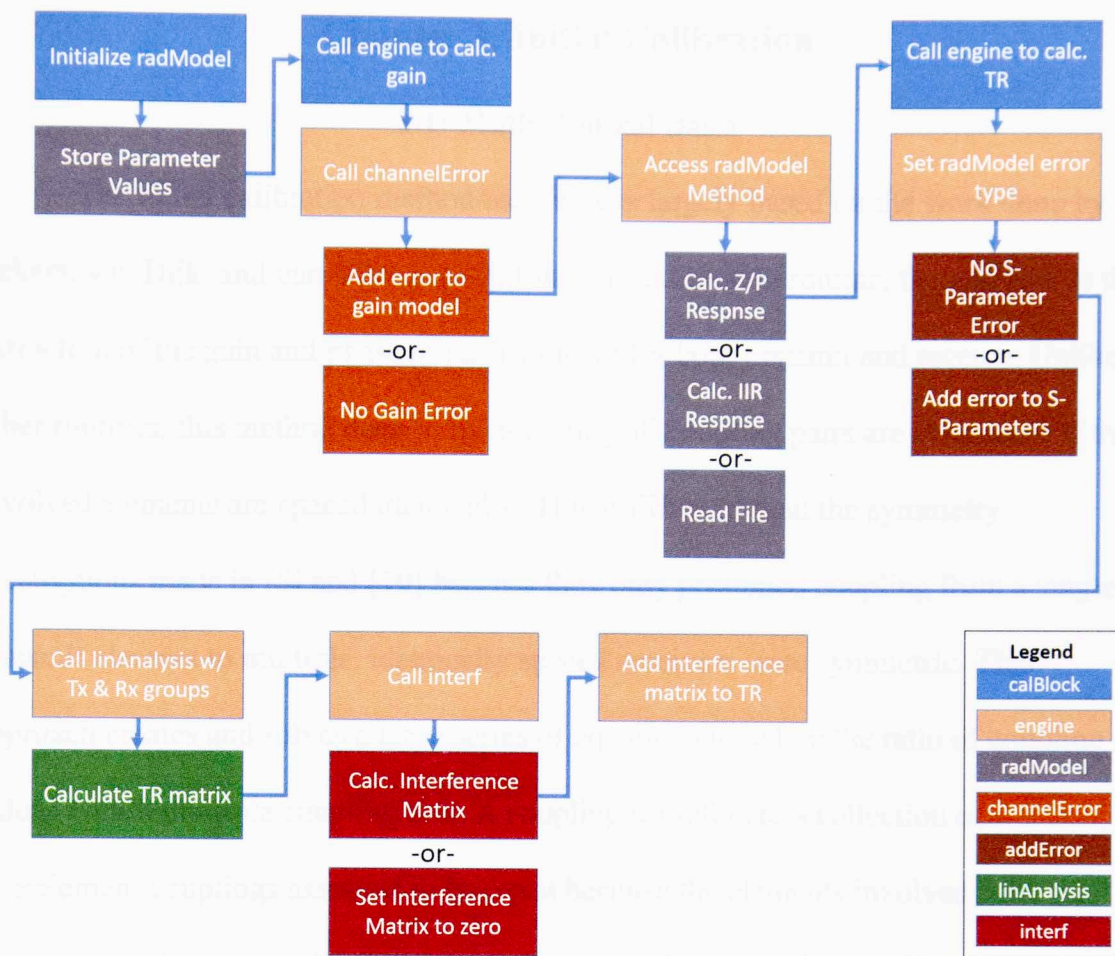


Fig. 3: Block diagram illustrating the sequential order of function calls.

On the top level, the calBlock is responsible for initializing radModel and calling the engine function under different scenarios. This is also where the calibration algorithms are executed and all user interaction takes place. Once prompted, the engine function calls the functions involved in calculating the gain, the TR matrix, and the interference matrix, returning the results to the calBlock for analysis by the user or the calibration methods. It is worth noting that the current methodology calls engine multiple times. The first few calls dictate how the gains should be calculated and how the errors should be applied to the gains and the s-parameters. The last call handles the generation of the TR and any interference matrices.

Chapter 3: Initial Calibration

3.1: Mathematical Basis

The initial calibration method used here is largely based on the work done by Bekers, van Dijk, and van Vliet in [12]. Like any calibration routine, the end goal is the extraction of the gain and phase of each element for both transmit and receive. Unlike other routines, this method does so by assuming all coupling pairs are symmetric if the involved elements are spaced identically. This is different than the symmetry assumptions made in [5] and [10] because they only presumed coupling from a single transmit element to multiple, identically spaced receivers were symmetric. This approach creates and solves a large series of equations based on the ratio of coupling values within multiple coupling sets. A coupling set refers to a collection of interelement couplings assumed to be equal because the elements involved have the same geometric spacing. For instance, in a rectangular array, the couplings between all elements and their neighbors one row directly above or below make up the [1,0] coupling set. [1,0] indicates the set's normalized row and column distances. The distances can be written generally as:

$$\left[\frac{D_{row}}{\Delta_{row}}, \frac{D_{col}}{\Delta_{col}} \right] = [q, p], \quad (8)$$

where D_{row} and D_{col} refer to the row and column distance, and Δ_{row} and Δ_{col} are the row and column spacing. In this work, the q and p definitions are flipped from the convention in [12], where q and p are column and row spacing respectively. This change was made to match the standard MATLAB® convention of listing the row first. Equations can be created by finding the ratio of pairs of elements within a set. So, if the

couplings $D_{xy}=T_xC_{yx}R_y$ and $D_{zw}=T_zC_{zw}R_w$, where T_n is the n^{th} complex transmit element channel gain, R_m is the m^{th} complex receive element channel gain, and C_{nm} is the term at the $[n,m]$ index from the previously defined \mathbf{C} matrix, the coupling ratio can be expressed as:

$$K_{xyzw} = \frac{D_{xy}}{D_{zw}} = \frac{T_x C_{xy} R_y}{T_z C_{zw} R_w}. \quad (9)$$

The C_{nm} terms, which represent the effective coupling, are assumed to be equal within a coupling set, so they will cancel out. If this assumption does not hold true, additional errors will appear. The coupling values are complex, so converting to the logarithmic domain allows us to separate the magnitude and phase portions of the coefficients as well as replace the multiplication and division with addition and subtraction. The resulting equations are:

$$\ln|T_x| + \ln|R_y| - \ln|T_z| - \ln|R_w| = \ln|D_{xy}| - \ln|D_{zw}|, \quad (10)$$

$$\angle T_x + \angle R_y - \angle T_z - \angle R_w = \angle D_{xy} - \angle D_{zw}. \quad (11)$$

Equations of this form can be written for every ratio in a coupling set and combined into two large matrix equations of the form $\mathbf{AE} = \mathbf{K}$. For the previous four element example, the matrix equations are written as:

$$\begin{bmatrix} 1 & 0 & -1 & 0 & 0 & 1 & 0 & -1 \\ & & \vdots & & & & & \end{bmatrix} \begin{bmatrix} \ln|T_x| \\ \ln|T_y| \\ \ln|T_z| \\ \ln|T_w| \\ \ln|R_x| \\ \ln|R_y| \\ \ln|R_z| \\ \ln|R_w| \end{bmatrix} = \begin{bmatrix} \ln|D_{xy}| - \ln|D_{zw}| \\ \vdots \end{bmatrix}, \quad (12)$$

$$\begin{bmatrix} 1 & 0 & -1 & 0 & 0 & 1 & 0 & -1 \\ & & \vdots & & & & & \end{bmatrix} \begin{bmatrix} \angle T_x \\ \angle T_y \\ \angle T_z \\ \angle T_w \\ \angle R_x \\ \angle R_y \\ \angle R_z \\ \angle R_w \end{bmatrix} = \begin{bmatrix} \angle D_{xy} - \angle D_{zw} \\ \vdots \end{bmatrix}. \quad (13)$$

The sparse matrix on the left, \mathbf{A} , does not depend on the measured coupling values, making it possible to calculate it ahead of time. Conveniently, it is the same for the magnitude and the phase equations, so only one matrix is necessary for both. The temptation is to find the pseudoinverse of \mathbf{A} , recreate \mathbf{K} from measurements, and multiply the two to retrieve transmit and receive coefficients. Unfortunately, the use of ratios to construct \mathbf{A} means it can never achieve full column rank, meaning $\mathbf{A}^{-1}\mathbf{A} \neq \mathbf{I}$, where \mathbf{I} is the identity matrix and \mathbf{A}^{-1} is the pseudoinverse of \mathbf{A} . Hence, it is necessary to treat one transmit element and one receive element as references and assume their coefficient values are known. The altered equation can now be written as $\mathbf{G}\mathbf{E}_1 + \mathbf{B}\mathbf{E}_2 = \mathbf{K}$. \mathbf{G} is a modified version of \mathbf{A} in which the columns corresponding to the reference elements have been removed. These two columns constitute the matrix \mathbf{B} . \mathbf{E} is also

split, with the reference elements in \mathbf{E}_2 and the remaining elements in \mathbf{E}_1 . If T_x and R_x are used as references, Equation 6 becomes:

$$\begin{bmatrix} 0 & -1 & 0 & 0 & 0 & -1 \\ & & \vdots & & & \\ & & & & & \end{bmatrix} \begin{bmatrix} \ln|T_y| \\ \ln|T_z| \\ \ln|T_w| \\ \ln|R_y| \\ \ln|R_z| \\ \ln|R_w| \end{bmatrix} + \begin{bmatrix} 1 & 1 \\ \vdots & \end{bmatrix} \begin{bmatrix} \ln|T_x| \\ \ln|R_x| \end{bmatrix} = \begin{bmatrix} \ln|D_{xy}| - \ln|D_{zw}| \\ \vdots \end{bmatrix}, \quad (14)$$

The phase version of Equation 8 is not shown, but it follows the same pattern demonstrated. This new, smaller matrix, \mathbf{G} , is full column rank, meaning $\mathbf{G}^{-1}\mathbf{G} = \mathbf{I}$, where \mathbf{G}^{-1} is the pseudoinverse of \mathbf{G} . In this text, \mathbf{G}^{-1} is referred to as the inverse coupling ratio matrix (ICRM). It is worth noting that $\mathbf{G}^{-1}\mathbf{B}$ results in negative ones such that all transmit and receive levels are relative to the references in \mathbf{E}_2 . Now, the transmit and receive coefficients in \mathbf{E}_1 can be recovered by multiplying \mathbf{G}^{-1} by \mathbf{K} , and the total list of coefficients can be assembled by converting back from the logarithmic domain and reintroducing the assumed reference values at the proper indexes.

As mentioned briefly before, calculation of the matrix \mathbf{A} does not depend on any measurements. This lends itself to breaking up the process into an online portion and an offline portion, as suggested in [12]. The offline portion consists of creating \mathbf{A} , reducing it to \mathbf{G} , and calculating \mathbf{G}^{-1} . During online operation, the coupling ratios in \mathbf{K} must be measured and multiplied by \mathbf{G}^{-1} to recover the desired coefficients. Calculating the pseudoinverse ahead of time significantly speeds up the calibration process, as the matrix \mathbf{G} will be very large, and finding its pseudoinverse will be time consuming.

3.2: Software Implementation

Implementing the method described above in software is a relatively straightforward matter. Predictably, the overall matrix to be inverted can be written as smaller matrices for each coupling set stacked with the last row of one above the first row of the next. As a result, the matrix for each set can be generated with a series of loops which steps through each element and stores the element pairs belonging to each coupling set. This list of pairs is used to create another array of all possible coupling ratios where the first two numbers represent the numerator's transmit and receive elements, and the second two numbers represent the denominator's transmit and receive elements. The array elements are denoted using a linear indexing scheme, so the columns of the large matrix can be identically numbered for the transmit side. The receive side follows the same pattern, but each index is offset by the total number of elements. The large matrix is then easily created by adding one to the columns representing the first two elements in each row of the ratio array and subtracting one from the columns representing the second two elements. This process is repeated for every desired coupling set, and the results are combined to generate the overall matrix.

Once the complete matrix is created, the columns for the reference transmit and receive must be removed, a simple operation in MATLAB®. This slightly reduced matrix is then inverted to create the ICRM using MATLAB®'s built in pseudo-inverse function.

The next step is calculating the ratios of the measured couplings. The coupling is approximated by the TR matrices returned when engine is called with a single element in transmit and receive groups. The right hand side of Equations 6 and 7 are then found

using the TR values. At this point, all but the reference channel coefficients can be found according to the last few steps in Section 3.1. Finally, the reference transmit and receive elements are reintroduced by inserting ones into the magnitude array and zeroes into the phase array to account for the assumed values of $1e^{j0}$.

All of this is done over a wide frequency bandwidth. The reference elements will have their own frequency responses, so their channel gains may not be $1e^{j0}$ at every frequency. As a result, assuming the reference elements have channel gains $1e^{j0}$ normalizes every element channel to the references' value at each frequency.

In software, each matrix is three dimensional, and the third dimension is treated as frequency. This results in several additional for loops in the code, but it is by far the most intuitive way of expanding the methods to a wide bandwidth.

3.3: Coupling Set Analysis

At the core of the method described in [12] is the assumption that all interelement couplings within a coupling set are equal. For instance, the coupling between elements 24 and 25 should be the same as between elements 57 and 58. Of course, this is not the case in practice due to things like edge effects and physical differences in the array. For that reason, it is important to examine and quantify exactly how similar the values within a coupling set are. Figure 4 contains overlapping plots of the frequency dependent coupling values within each set, visually demonstrating the likeness of the sets' couplings. One can see qualitatively that almost every coupling set diverges in at least one frequency range and that the sets differ most often at higher frequencies for this particular antenna. This illustrates why it is important to be selective when choosing the coupling sets used to generate the inverse matrix. Figure 5 serves as

a way to quantify these differences within sets. It plots the standard deviation of each group across frequency, showing the greater deviations at high frequencies. However, the equations for initial calibration rely more heavily on the ratios between coupling values than the values themselves. The standard deviations of the ratios are plotted in Figure 6, and Figure 7 contains the overlapping plots of the ratios within each set.

Certain sets were left out of Figure 6 due to their unsuitability. These correspond to the plots in Figure 7 which diverge greatly at high frequencies. Their standard deviations were at least two times greater than those shown in Figure 6 at high frequencies, so including them in the analysis would have made it more difficult to recover the elements' frequency response in that range.



Figure 7: Overlapping plots of the coupling values (magnitude in dB) within each set across frequency.

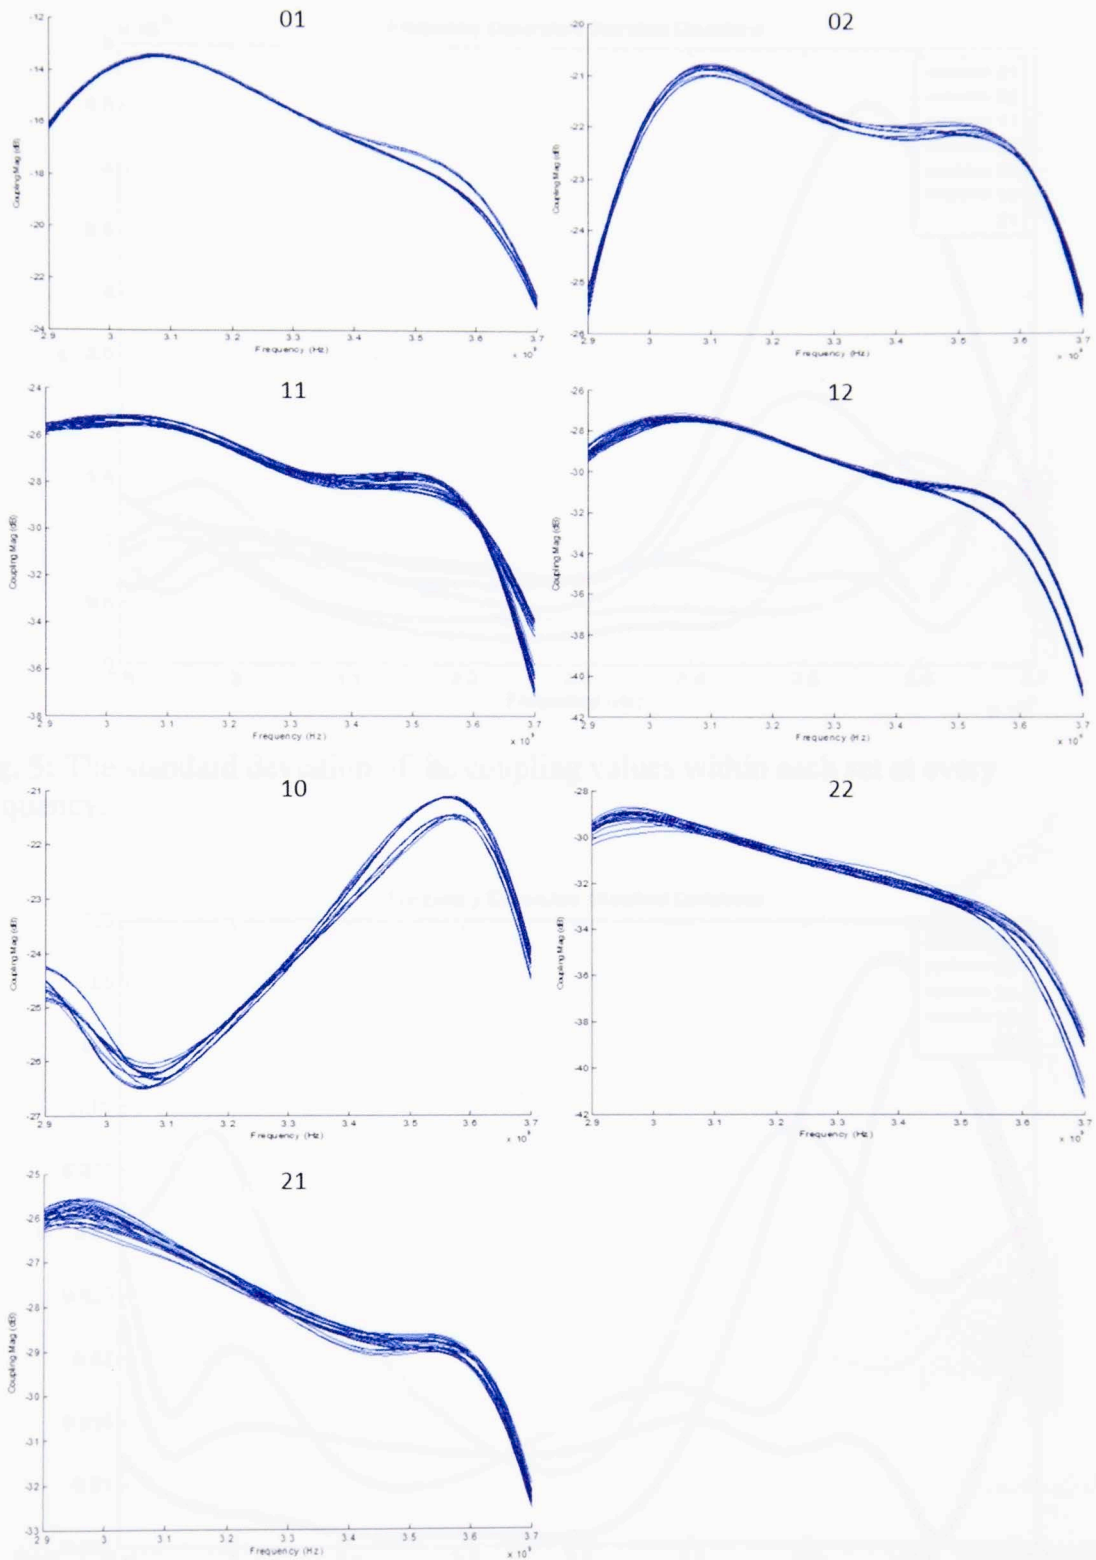


Fig. 4: Overlapping plots of the coupling values (magnitude in dB) within each set at every frequency.

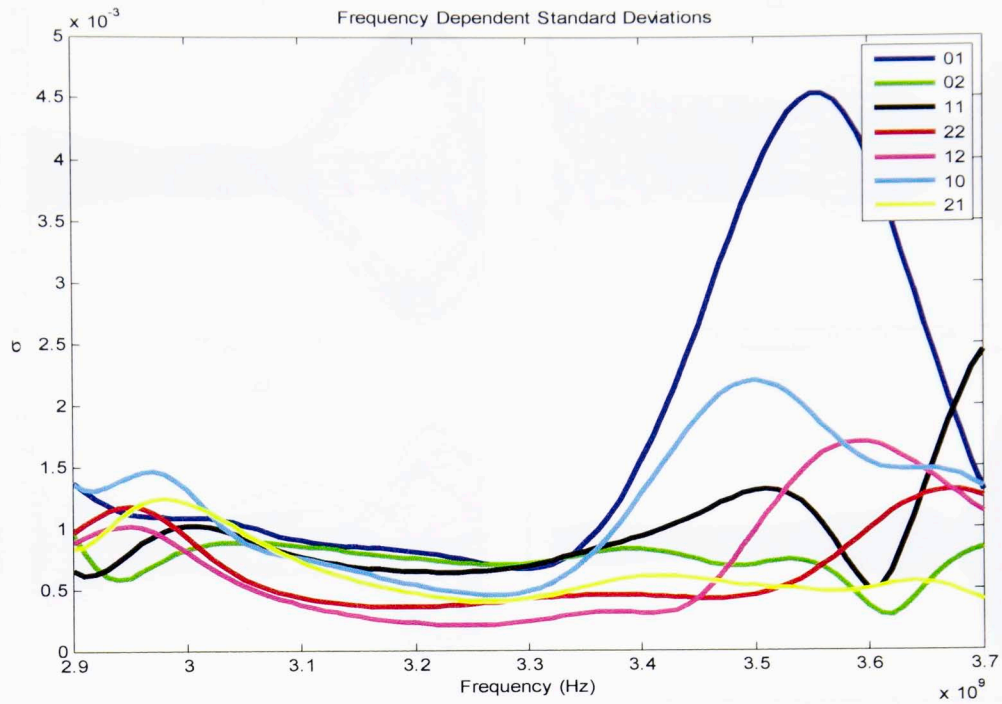


Fig. 5: The standard deviation of the coupling values within each set at every frequency.

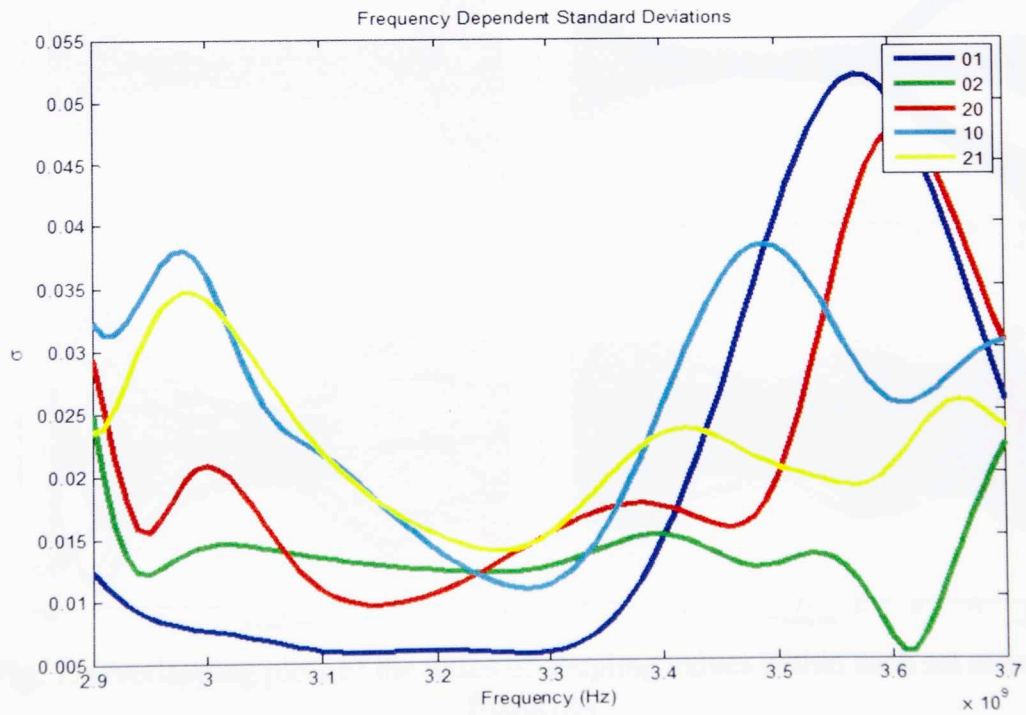


Fig. 6: The standard deviation of the ratios of coupling values within each set across frequency.

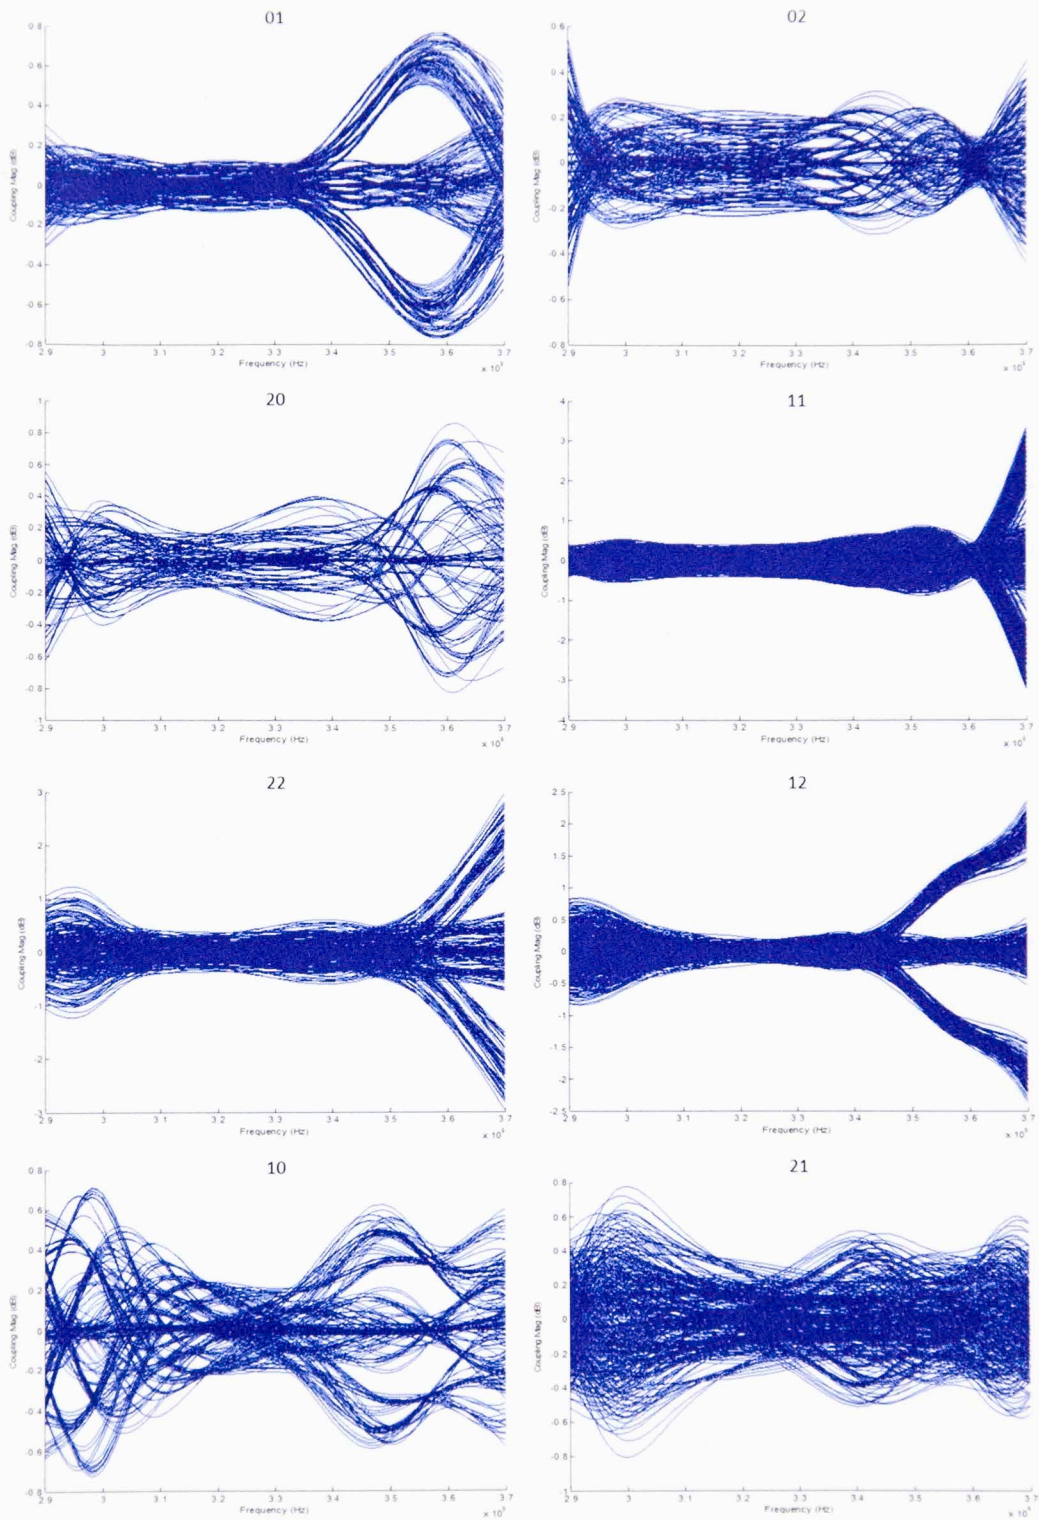


Fig. 7: Overlapping plots of the ratios of coupling values within each set at every frequency

3.4: Results

To test the initial calibration approach, a five element by five element subset of a nine by nine array was used, and the center element was picked as the reference. The [1,0], [0,1], [0,2], and [2,1] coupling sets were used to reach the appropriate rank. No distance criterion were enforced to prune down the coupling sets. A complex error with a standard deviation of .01 was added to the coefficients of the IIR gain model. While only the transmit results will be presented here, the receive chains exhibit all the same behaviors, so showing them would be redundant. First, the s-parameters were changed to an array of all ones to ensure the method worked in an ideal situation. While ones may not be a physically realistic s-parameter value, this setup is valid because the coupling values are being used in ratios. When the s-parameters are identical, as they are in this case, dividing the coupling values will cancel out the matching s-parameter terms. As a result, the s-parameter matrix can be filled with any value and it will be an ideal simulation as long as all of the values are the same. The results, seen in Figure 8, demonstrate that the estimate perfectly tracks the response across frequency. It tracked so closely that a .005 offset had to be added if both traces were to be visible.




Fig. 8: Frequency response of element 23's transmit chain when ideal s-parameters are used. A .005 offset has been added for the sake of visibility.

However, the accuracy of the results are dependent on how close the reference element is to ideal. Since the gains are all calculated as a ratio with the reference element, the errors in the reference element are spread across the other elements. Figure 9 shows what happens when random error is added exclusively to the reference element. The plot on the left corresponds to the reference element's response, and the plot on the right represents all others. The red traces are the estimated channel responses produced using the initial calibration routine, and the blue traces are the ratio of the altered gains to the center element's nominal value. In other words, the blue traces are what one would expect the channel responses to be given errors were only added to the center

element. However, forming a ratio with that altered center element essentially swaps the response behaviors, resulting in errors appearing on the unaltered channels and an apparently ideal reference, as demonstrated by the red traces in Figure 9. Initial calibration must be done without extensive prior knowledge of the array, so, in practice, there would be no way of identifying the errors came from the reference element and not the other channels.




Fig. 9: (Left) The altered center reference element response normalized to itself, in red, and normalized to the nominal reference response, in blue, (Right) All other channel responses relative to the altered reference, in red, and relative to the nominal reference channel, in blue

Of course, ideal s-parameters can make any result look better, so the method had to be tested with more realistic parameters as well. For these tests, the nine by nine array described in Section 1.3 and stored in the radModel object was used. Again, the center element was assumed to be an ideal reference, and error was added to every other channel. Element 43's frequency response is shown in Figure 10. The results are still a reasonable estimation of the frequency response, but Figure 10 is indicative of how the method can have trouble tracking the response at higher frequencies. As discussed in

Section 3.3, this is likely due to the higher deviation in the coupling sets at high frequencies. The linchpin of this approach is the assumption that all values within a coupling set are equal, and if the standard deviations grow too large, the accuracy of the results will suffer.

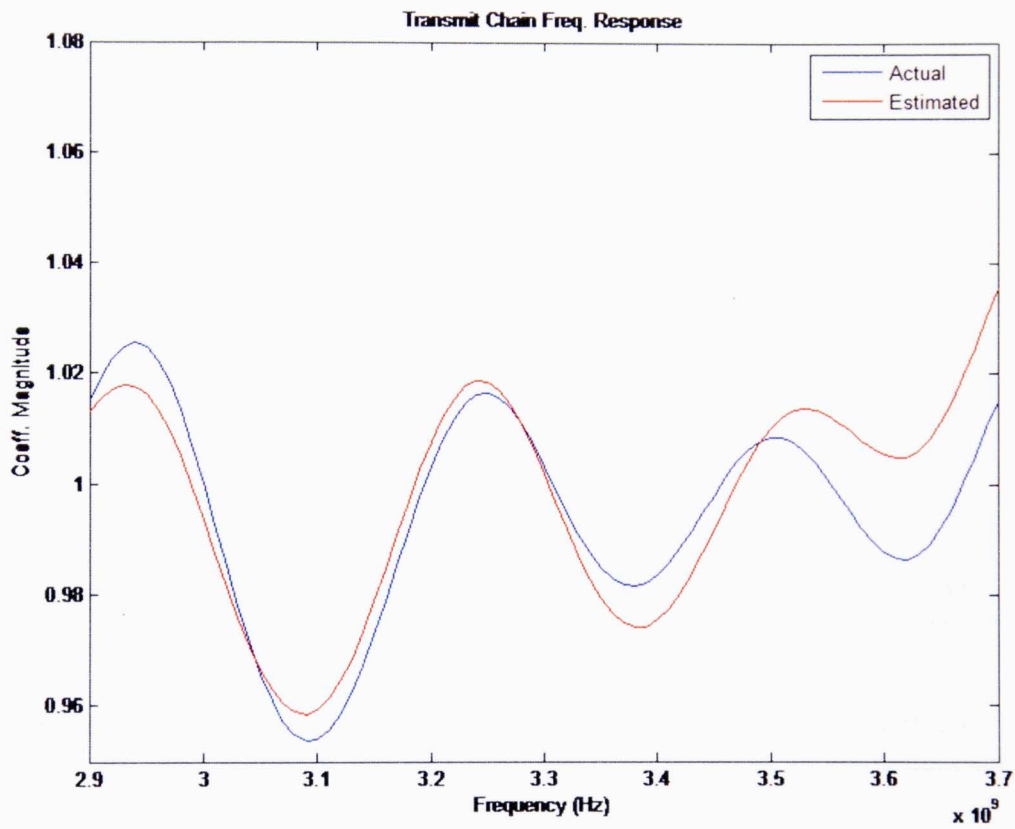


Fig. 10: Frequency response of element 43's transmit chain when simulated s-parameters are used. Random complex error has been added to all elements except the center element, which is used as the reference.

Figures 11 and 12 show the responses for elements 21 and 57 respectively. These two plots illustrate the overall range in accuracy this method produces. The majority of the responses are very similar to element 21 and element 43's precision, and element 57 is likely the worst approximation from the whole array. However, even 57's

estimated response roughly tracks the shape of the response, so equalization is not a lost cause.




Fig. 11: Frequency response of element 21's transmit chain when simulated s-parameters are used. Random complex error has been added to all elements except the center element, which is used as the reference.

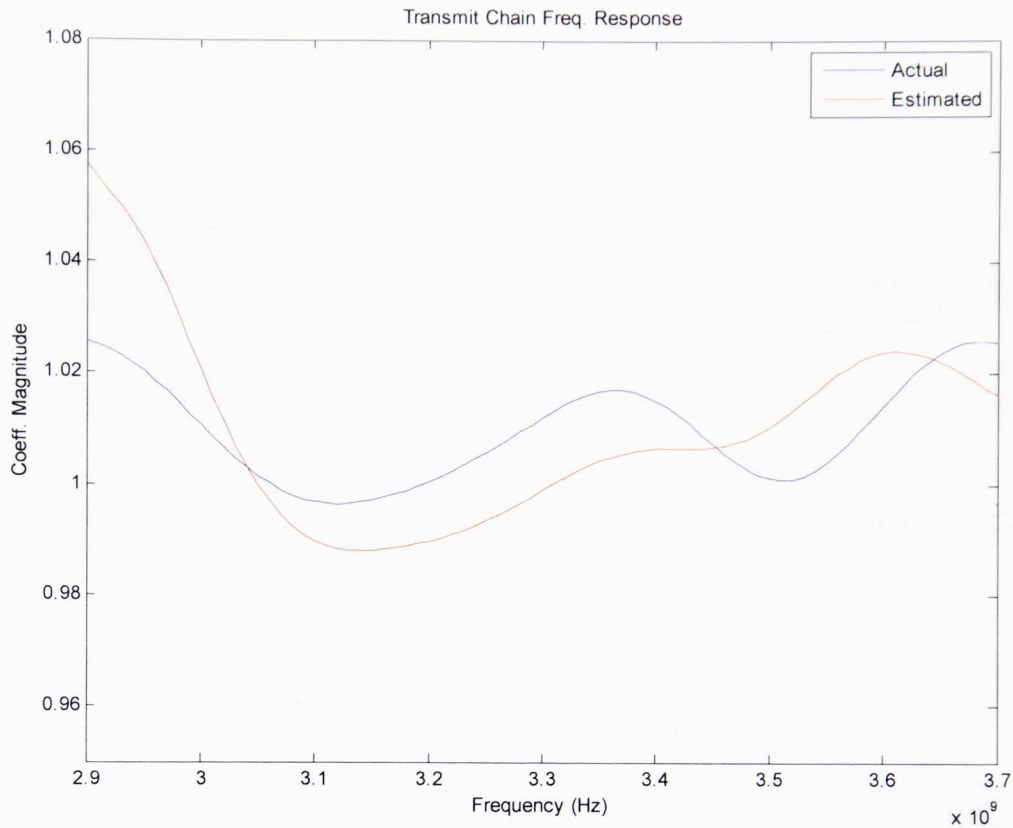


Fig. 12: Frequency response of element 57's transmit chain when simulated s-parameters are used. Random complex error has been added to all elements except the center element, which is used as the reference.

Another way of judging the calibration method's efficacy is by investigating how it tracks gains across the array face. Figures 13 and 14 show what happens when the magnitude and phase of a single element is changed manually. The magnitude of the transmit gain coefficients at 3.17 GHz are plotted in Figure 13 as a surface to demonstrate how closely the method tracks amplitude changes. Figure 14 does the same for the phases. The approximation never differs by more than .02 in magnitude and 2 degrees in phase. For the results in Figure 15, error was again only added to a single element, but this time it was generated by randomly changing the IIR coefficients for a

single channel. As was the case with the prior test, the magnitude estimate never differs from the actual value by more than .02.




Fig. 13: Results at 3.17 GHz when a single element is manually altered, (left) Estimated element gain magnitudes, (center) Actual element gain magnitudes, (right) Magnitude of the difference between estimated and actual gain magnitudes.




Fig. 14: Results at 3.17 GHz when a single element is manually altered, (left) Estimated element gain phases at 3.17 GHz, (center) Actual element gain phases, (right) Magnitude of the difference between estimated and actual gain phases




Fig. 15: Results at 3.09 GHz when a single element is randomly altered, (left) Estimated element gain magnitudes, (center) Actual element gain magnitudes, (right) Magnitude of the difference between estimated and actual gain magnitudes

Figures 16 and 17 show the results of what is by far the most realistic test case, in which random error with a standard deviation of .05 is applied to the IIR numerator coefficients of every channel, and random error with a standard deviation of .0167 is applied to the channels' poles. These errors were applied with a newer revision of the channelError function to more closely resemble actual value changes. Even under these conditions, the estimates and the actual values are still similar in both a qualitative and a quantitative sense. The magnitude is never off by more than .025, and the phase estimate is always within 1.5 degrees of the actual value. Additionally, the surfaces actually look alike; the estimate captures the trends and allows the user to visually judge their accuracy.




Fig. 16: Results at 3.3 GHz when elements are randomly altered, (left) Estimated element gain magnitudes, (center) Actual element gain magnitudes, (right) Magnitude of the difference between estimated and actual gain magnitudes




Fig. 17: Results at 3.3 GHz when all but the center element are randomly altered, (left) Estimated element gain phases, (center) Actual element gain phases, (right) Magnitude of the difference between estimated and actual gain phases

For the sake of comparison, results from the last test were plotted on an identical scale to that used in results from [12]. The relative amplitude and phase difference plots from [12] are plotted alongside the results calculated using the reverse engineered method in Figure 18. The tests do differ in that [12] worked with an eight element by eight element array rather than a five by five, but the errors of the calibration done here fall safely within the accuracy range the original work reported.

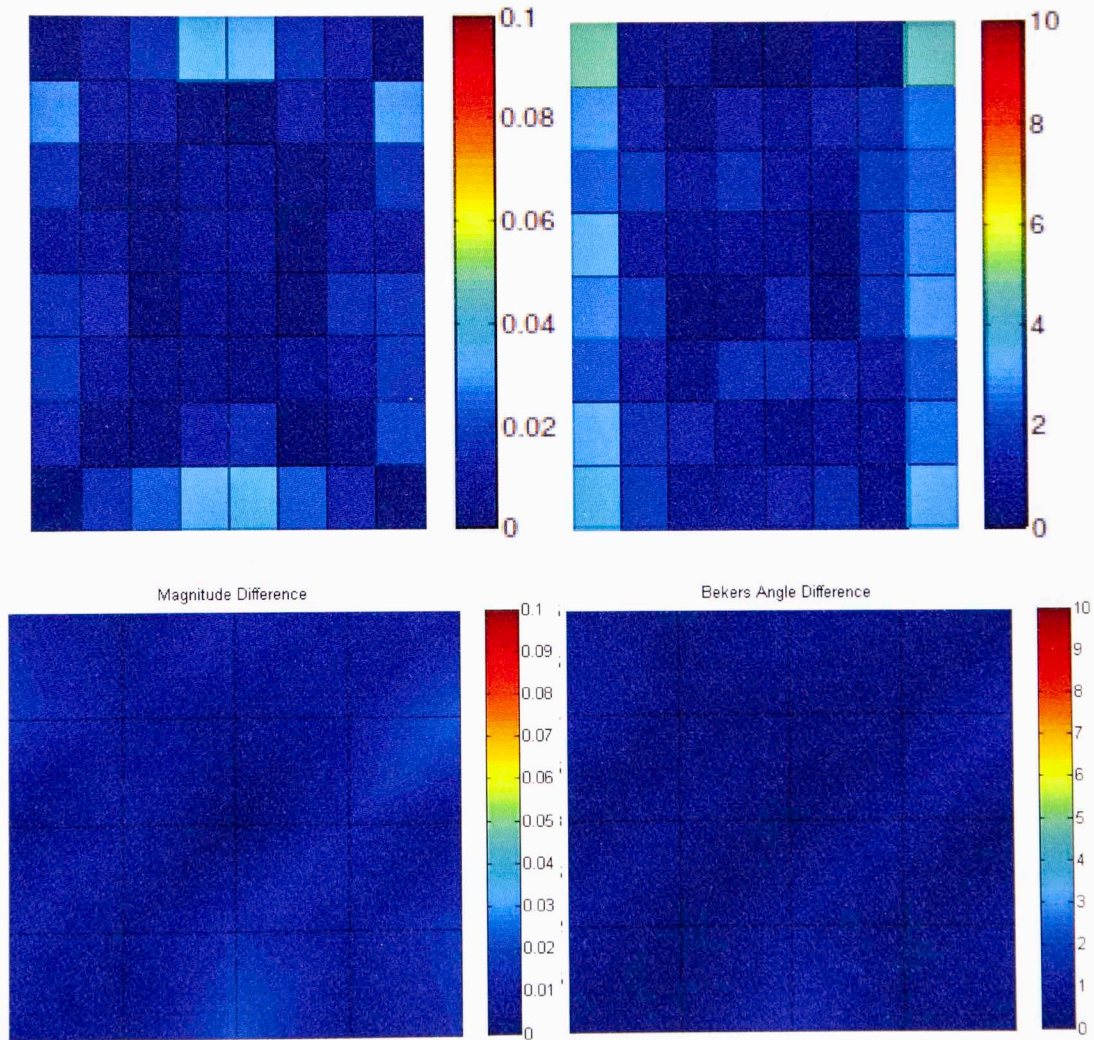


Fig. 18: A comparison of the original absolute magnitude and phase differences from [12] with those calculated using the reverse engineered methodology

As mentioned previously, the efficacy of the initial calibration routine will depend on the deviation within each coupling set. One practical consideration is then

the impact of edge effects on calibration. Edge effects are extremely likely to disrupt the symmetry assumptions. To test this idea, the five by five subset was relocated to the bottom right corner of the array, and the calibration routine was run to recover random errors applied to each element with the latest channelError function. Figures 19 through 22 plot the differences between estimated amplitudes and phases for each element as a surface. The errors at both 3.3 GHz and 3.55 GHz are shown. These plots illustrate the impact the edge effects have on the estimation. The corner subset had higher error levels for both amplitude and phase, especially at higher frequencies. From this analysis, it can be concluded the array edges must be taken into account when choosing to implement this initial calibration routine.

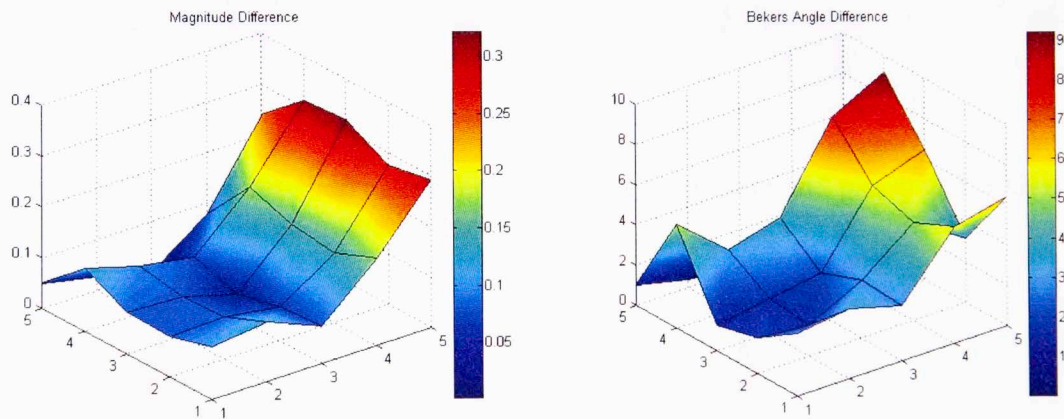


Fig. 19: Differences between estimated and actual amplitude and phase at 3.3 GHz when the subset is placed in the corner




Fig. 20: Differences between estimated and actual amplitude and phase at 3.55 GHz when the subset is placed in the corner




Fig. 21: Differences between estimated and actual amplitude and phase at 3.3 GHz when the subset is placed in the center




Fig. 22: Differences between estimated and actual amplitude and phase at 3.55 GHz when the subset is placed in the center

Chapter 4: In-Situ Calibration

4.1: Mathematical Basis

If a baseline TR matrix can be generated to be used as a “golden standard”, the method proposed in [13] can be used to evaluate the errors of subsequent measurements. If the subsequent matrix is \mathbf{TR}' , and the baseline matrix is \mathbf{TR} , then the normalized coupling matrix, \mathbf{K} , is defined as

$$\mathbf{K} = \mathbf{TR}' \circ \mathbf{TR}, \quad (15)$$

where the \circ symbol denotes the inverse Hadamard product. If the s-parameters and reflection coefficients remain constant, \mathbf{K} can be written as an outer product of the error vectors. As a result, a singular value decomposition (SVD) can easily recover the errors on transmit and receive up to an arbitrary amplitude and phase reference between them [13]. The SVD will yield a left singular vector matrix \mathbf{U} , a singular value vector \mathbf{S} , and a right singular vector matrix \mathbf{V} such that

$$\mathbf{K} = \mathbf{USV}', \quad (16)$$

with \mathbf{V}' representing the conjugate transpose of \mathbf{V} . \mathbf{U} and \mathbf{V}' are given as normalized matrices, so they must be de-normalized by multiplying \mathbf{U} by the square root of the number of receivers and multiplying \mathbf{V}' by the square root of the number of transmitters. \mathbf{S} must be divided by the two multiplication factors as well to maintain equality. The errors can then be split evenly over all elements by multiplying both \mathbf{U} and \mathbf{V}' by the square root of the dominant singular value. The transmit errors and receive errors are then easily extracted from \mathbf{V}' and \mathbf{U} respectively. Alternatively, the values in \mathbf{U} and \mathbf{V}' can be set relative to a reference element in the transmit or receive groups and all the errors can be spread over the other group by multiplying the

corresponding matrix by the dominant singular value. Both of these approaches are used in the tests described in Section 4.4. A more detailed explanation of the mathematics can be found in [13].

4.2: Combining Multiple Iterations

In order calculate transmit and receive errors for an entire array, multiple calculations with different transmit and receive groups must be performed. This necessitates a means of connecting the multiple results. For this work, the five different transmit/receive schemes in Figure 23 were used. The first four groups were chosen because each arrangement has transmit and receive elements which overlap with the previous and the next calculation. However, the first four configurations do not include the center element as a transmitter, so the special fifth case must be added.

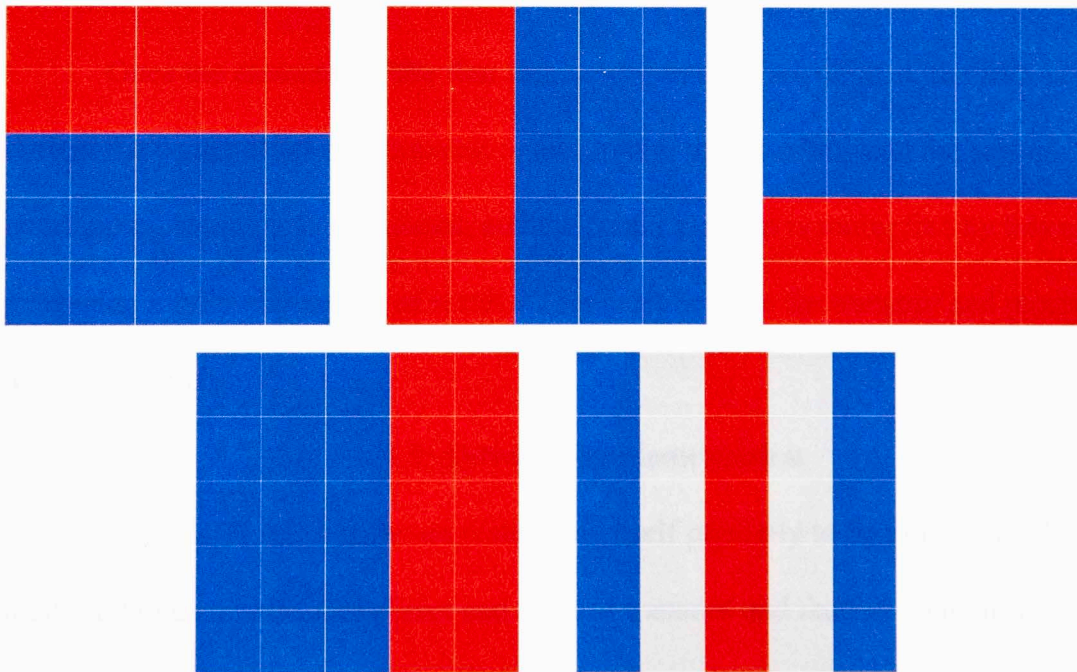


Fig. 23: The five transmit/receive configurations used to calculate the complete array error. The red represents transmitting elements, the blue represents receiving elements, and gray elements are unused

Assume R_1 and R_2 are vectors containing the complex values for the overlapping elements within the region in the first and second measurements, respectively. If C_{12} is the mean of the element-wise ratios of R_1 to R_2 , then the entire region of the second measurement can be aligned with the first by multiplying each element by C_{12} . This idea cascades, so the third measurement can be aligned by multiplying by both C_{12} and C_{23} and so on. However, there are four different possible multiplication factors for each measurement depending on which of the four configurations is measured first. To account for this, the four factors are averaged together, and the elements of interest are multiplied by that mean. It is worth noting that $C_{12}C_{23}C_{34}C_{41}=1+\varepsilon$, where ε is the error. As an example, the average factor, for the first arrangement in Figure 23 will be

$$\frac{1 + C_{41} + C_{34}C_{41} + C_{23}C_{34}C_{41}}{4} \quad (17)$$

Once the elements of the first four configurations are aligned, the fifth is used to correct the center element's transmit value. Again, the ratio between the sets of overlapping elements is calculated, and the center element is multiplied by this ratio, producing a fully aligned set of results. This must be done for transmit and receive levels separately.

4.3: Software Implementation

The MATLAB® software suite lends itself perfectly to this in-situ calibration method. Given a radModel object and a list of transmit and receive element indexes, the program uses the engine function to call linAnalysis to generate TR matrices for both the nominal and altered versions of the array. The inverse Hadamard product is calculated using MATLAB®'s element-wise division operator, and the built-in SVD function generates \mathbf{U} , \mathbf{S} , and \mathbf{V} . The de-normalization and distribution of errors is then

performed using basic operators, and transmit and receive values are extracted from \mathbf{V} and \mathbf{U} . Each step is repeated multiple times within for loops to calculate the values at multiple frequency points. As before, the third dimension is used to represent the different frequencies.

If the transmit and receive groups from Figure 18 are being aligned, the overlapping element indexes must be declared as well. The method for calculating \mathbf{V} and \mathbf{U} is then repeated for each configuration, and the ratios are calculated and averaged to create the scale factors for the first four configurations. \mathbf{V} and \mathbf{U} are calculated once more for the fifth element arrangement, and it is brought in line with the rest of the array.

4.4: Results

The in-situ calibration method was tested on the same five by five subset of a nine by nine array as the initial calibration. The transmit group consisted of the elements in the first two columns, elements 21-25 and 30-34, and the receive group contained the last three columns, elements 39-43, 48-52, and 57-61. All tests were performed using the same simulated s-parameters as well. A complex random Gaussian error with a standard deviation of .01 was applied to the IIR filter coefficients modeling the channel gains.




Fig. 24: (left) Estimated frequency response of element 60's receive chain in blue and the actual response in red, (right) Overlapping plots of the difference between the estimated and actual responses for every element in the receive group




Fig. 25: (left) Estimated frequency response of element 34's transmit chain in blue and the actual response in red, (right) Overlapping plots of the difference between the estimated and actual responses for every element in the transmit group

For the first test case, the errors were assumed to be split evenly between all elements. Figures 24 and 24 show the results for an element in the receive group and the transmit group respectively. For both transmit and receive, the estimates' magnitudes are never off by more than .009 within the entire frequency range.

The second test case demonstrates how the frequency response can be perfectly recovered if an ideal reference element is present to normalize the singular value decomposition results. The frequency response of a single transmit channel is plotted in Figure 26, but all transmit and receive channels exhibited the same ideal performance. The errors in the right plot of Figure 26 are so small they are likely a result of minor mathematical errors inherent in the software.

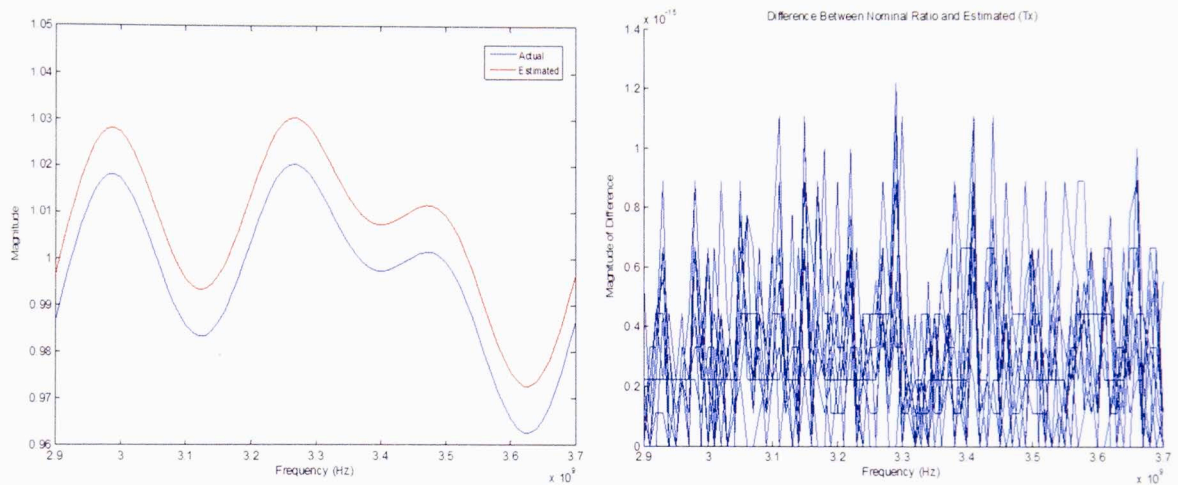


Fig. 26: (left) Estimated frequency response of element 34's transmit chain in blue and the actual response with a .01 offset in red, (right) Overlapping plots of the difference between the estimated and actual responses for every element in the transmit group

Unfortunately, the second test case is unrealistic in practice because it is unlikely a perfect reference element can be maintained, especially after an array has been fielded. In light of this, the third test case was identical to the second, but error was added to the reference element as well. The results for the receive chains of the 60th and the 49th elements can be seen in Figure 27. The estimate of element 60's response still tracks the actual frequency response without differing by more than a few hundredths. However, the same cannot be said of the estimate of element 49's response. In fact, it was perhaps the worst estimate of the response shape out of the whole receive group.

Figure 28 makes these same comparisons for the transmit group; the estimate of element 34's response is accurate to within a few hundredths, and the estimate of element 22's response is the worst of the transmit group estimates. The overall accuracy of this test case can be seen in Figure 29. Across the entire frequency range, the magnitude of the estimates never differ from the actual responses by more than .025. On its own, this is not a large error, but it is significantly more than the .009 difference produced when the errors are averaged across all elements.




Fig. 27: (left) Estimated frequency response of element 60's receive chain in blue and the actual response in red, (right) Estimated frequency response of element 49's receive chain in blue and the actual response in red




Fig. 28: (left) Estimated frequency response of element 34's transmit chain in blue and the actual response in red, (right) Estimated frequency response of element 22's transmit chain in blue and the actual response in red

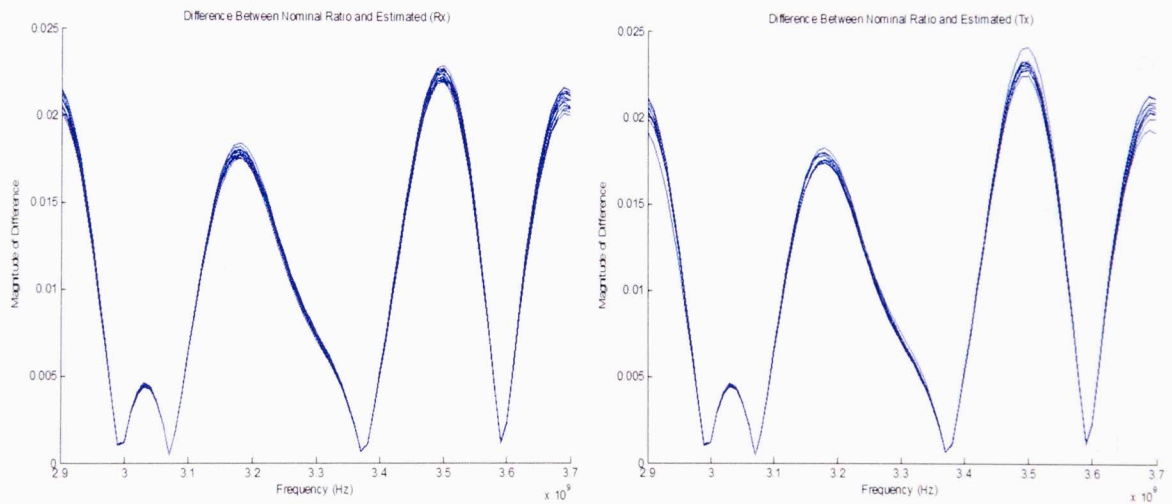


Fig. 29: (left) Overlapping plots of the difference between the estimated and actual responses for every element in the receive group, (right) Overlapping plots of the difference between the estimated and actual responses for every element in the transmit group

The choice between averaging errors and using a reference element depends largely on the nature of the reference element. If one can safely assume an element performs ideally, normalizing to that element is preferable. On the other hand, splitting the errors evenly over all elements will consistently produce accurate estimates without relying on the performance of a single channel.

4.5: Tracking S-Parameter Errors

Because this method assumes the s-parameters and reflections remain constant, it is important to have a means of checking the validity of this assumption. One way of doing so is by tracking the strength of the non-dominant singular values relative to the dominant one [13]. The normalized singular value sum, α , is used as a measure of this strength and is defined as

$$\alpha = \frac{1}{s_1} \sum_{n=2}^{\min(M,N)} s_n. \quad (18)$$

The magnitude of α can then be used to judge if the normalized coupling matrix behaves as an outer product of error vectors [13], which is only the case if the s-parameters and reflections remain the same. Ergo, the larger the value of α , the greater the changes in s-parameters or reflection coefficients.

To test the validity of this measure, the same five by five subset from Section 4.3 was used with the first two columns transmitting and the last three columns receiving. Errors from the SVD were split evenly amongst all elements and an additional complex error was added to the simulated s-parameters. α was calculated for 100 different error scale factors, ranging between 10^{-4} and 1. The results are shown in Figure 25. It can be easily seen that the magnitude of α tracks the error scale. To frame the result in terms of a measure one would have access to in practice, α was calculated again, but this time it was plotted as a function of the RMS element error. The RMS element error is defined here as

$$RMS = \sqrt{\frac{\sum_n |a_n - a_{n,l}|^2}{N}}, \quad (19)$$

if a_n is the n^{th} element's estimated complex gain, $a_{n,l}$ is the actual complex gain of the n^{th} element, and N is the total number of elements. The resulting plot can be seen in Figure 26. Once again, the α value accurately scales with the magnitude of the element errors. Both these results prove α 's worth as an error indicator; it could be tracked across measurements, and if it falls above a specified threshold, the results can be treated as invalid.

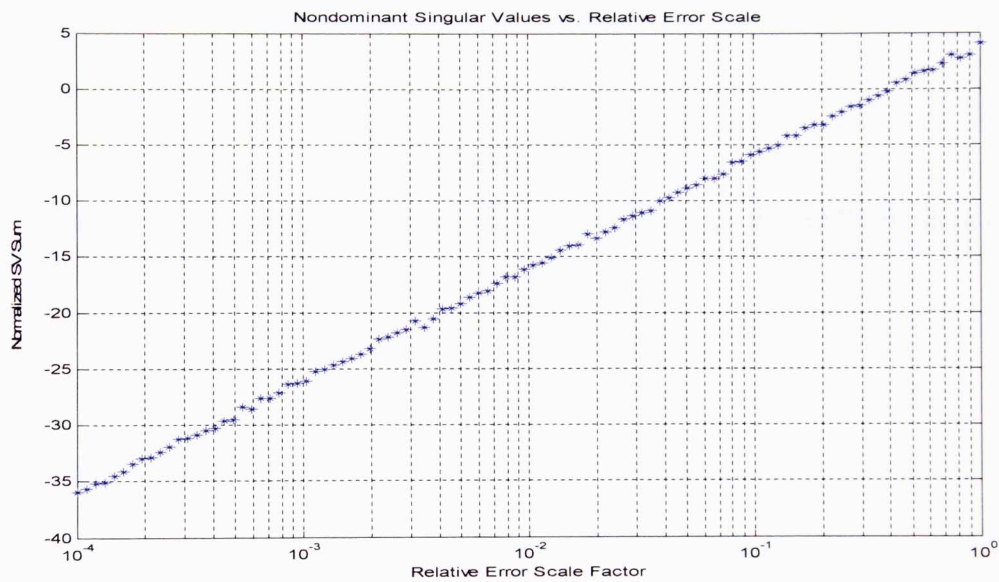


Fig. 30: The normalized singular value sum, α , plotted versus the error scale factor. Plot shown for 3.3

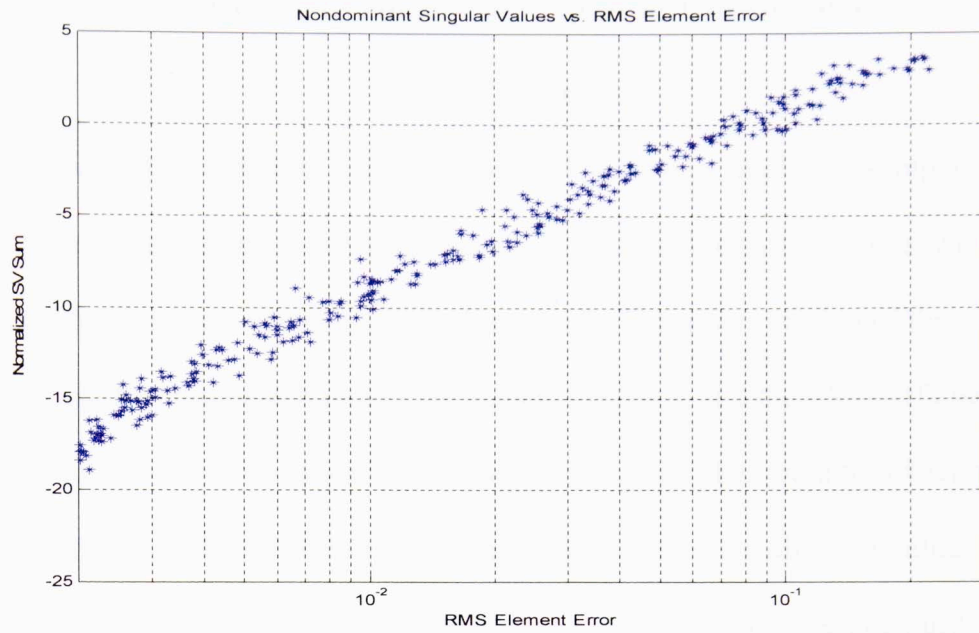


Fig. 31: The normalized singular value sum, α , plotted versus the RMS element error. Plot shown for 3.3 GHz.

Chapter 5: Calibration Method Comparison

5.1: Test Setup

Since both calibration methods produced promising results, it is important to do a direct numerical comparison between them. The same five by five subset with simulated s-parameters was used for both methods. The newer version of the channelError function which moves the zeroes, the poles, and the IIR coefficients every time and does so more realistically was used. Both approaches were tested with the standard deviation of the IIR numerator coefficients set to five different values between .001 and .063 and the standard deviation for the poles assumed to be one third of the numerator deviation every time. The s-parameter error scale factor was also set to five different values between .0001 and .2. The initial calibration method used the [1,0], the [0,1], the [0,2], and the [2,1] coupling sets to generate the ICRM and assumed the center element was the reference. The in-situ calibration used the five transmit and receive configurations in Figure 18, and the errors were split evenly across all elements.

5.2: Results

The RMS element errors, averaged for all frequencies, from both analyses were plotted on surface plots. The x and y axes represent the standard deviation of the coefficients and the scale factor for the s-parameter error, and the RMS element error is plotted on the z axis. Figure 32 shows the results from the initial calibration approximation, and Figure 33 shows the results using the in-situ calibration method.

The in-situ calibration produces less error than the initial calibration, especially as the error scale increases. In addition, the in-situ approximation is only effected by the gain errors as the s-parameter errors reach their highest levels. This is to be expected

because the in-situ calibration relies on a golden standard, so the s-parameter errors and gain errors are easily delineated. However, if the quality of the golden standard suffers, the method becomes more prone to error.

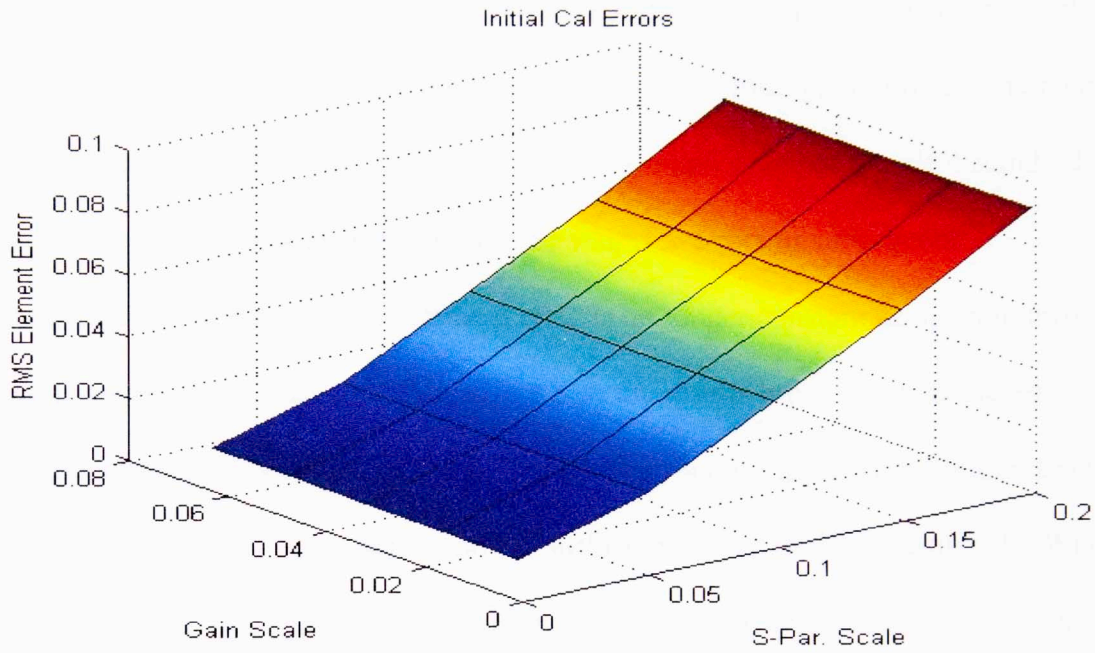


Fig. 32: RMS Element Error from the initial calibration routine as a function of the error scales.




Fig. 33: RMS Element Error from the in-situ calibration routine as a function of the error scales.

To better understand the limits of the in-situ calibration method, the previous test was repeated, but the ideal standard was replaced with a copy whose nominal s-parameters had complex random Gaussian error added. Complex error scales of .1, .2, .3, and .4 were tested. The calculations were repeated five times for each error scale, and the resulting RMS errors were averaged. Figures 34 through 37 present the resulting error surfaces. The first two results still retain the rough shape of the ideal standard's results, but there is now error even if the s-parameter errors of subsequent measurements are zero, and the element error is no longer constant as the gain error scale changes. Figures 36 and 37 illustrate how the results break down as the error in the golden standard becomes too high. The RMS errors in Figure 31 are higher than the initial calibration's at every point, and the surface becomes far less predictable. When an error scale of .4 is applied to the golden standard, the RMS errors skyrocket, as demonstrated in Figure 32. The averaging of multiple runs also hides the fact that a single result can have error peaks even greater than those in Figure 32, which could prove important if multiple measurements cannot be made in practice. Figure 38 is a plot of the 59th element's frequency response with an error scale of .4 applied to the standard and an error scale of .2 applied to the subsequent s-parameters; it serves as an example of how poorly in-situ calibration routine estimates the response in the presence of large s-parameter errors even if there are no gain errors. Fortunately, the value of α reflects these errors, and the initial calibration is unaffected by the changes in the standard. As a result, an α above a specified threshold could be used to signal a switch to the initial calibration routine to reduce errors.




Fig. 34: In-situ calibration errors when the golden standard has an error scale of .1 applied




Fig. 35: In-situ calibration errors when the golden standard has an error scale of .2 applied




Fig. 36: In-situ calibration errors when the golden standard has an error scale of .3 applied

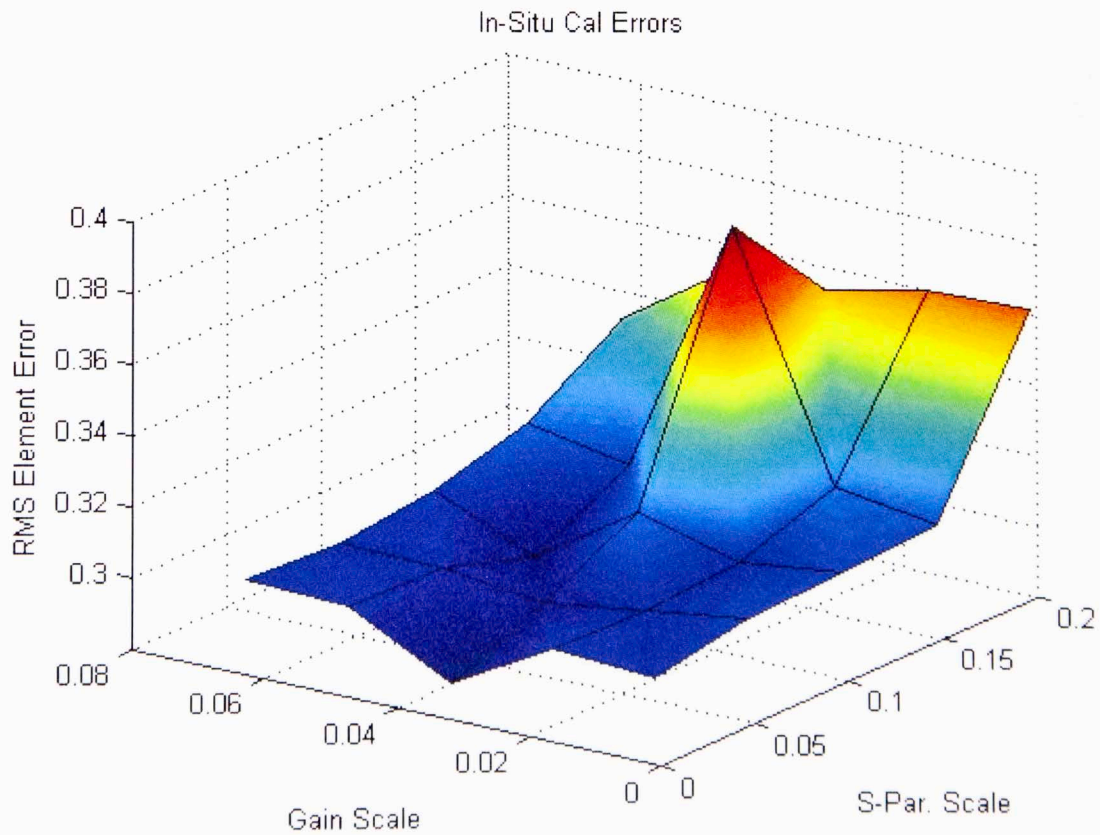


Fig. 37: Average of in-situ calibration errors from ten runs when the golden standard has an error scale of .4 applied




Fig. 38: Element 59's transmit channel frequency response when the golden standard has an error scale of .4, the measurement's s-parameters have an error scale of .2, and there is no gain error

Chapter 6: Conclusion

Based on the results in Chapters 4 and 5, it can be concluded that both methods of calibration presented provide a viable means of approximating gain and phase errors as they arise in a phased array radar. Under typical operating conditions, the “golden standard” method of in-situ calibration produces far less element error than the initial calibration because it incorporates an ideal model to more accurately identify errors. In addition, the RMS element error does not depend on the standard deviation of the gain error. However, the generation of an ideal model requires access to high fidelity equipment and a controlled test environment like an anechoic chamber or compact range. There are many situations in which utilizing this kind of test equipment is not feasible, and the generation of a high quality standard may not be possible. Additionally, the assumption the s-parameters will not change between measurements is less likely to hold in applications that involve extreme conditions. For instance, it is very likely the s-parameters of a space-based system would change from those measured on earth during the launch, so the creation of an accurate baseline may not be feasible.

The initial calibration procedure is a good alternative when access to a golden standard is limited. The ability to calculate the ICRM ahead of time can significantly reduce the processing time required for each calibration, producing substantial time gains over the life of the array. It is also possible to produce more accurate results than those shown here. The estimations relied on simulated coupling sets with relatively high standard deviations. As the deviations are reduced through careful design, the element

error will fall as well. Some form of iterative design with distance criteria implemented could also improve the initial calibration accuracy.

Overall, the calibration software framework presented here will be an excellent foundation for future calibration testing. More work could identify the mathematically optimal way to tie multiple in-situ measurements together. The radar model can be expanded to include radiation patterns and direction dependent effects for each antenna element. Currently, the radiating properties of the antennas are being ignored; the channels feeding the elements have been the main focus. However, to begin synthesizing a complete radiation pattern for the array, the angular dependence of the elements' transmitted or received signals must be accounted for. Once this has been addressed, the effects of calibration on beamforming can be examined as well. It is also possible to investigate how the calibration methods can be used on systems with different layouts such as cylindrical phased arrays. Automation of the calibration process would also be an interesting field to examine. More quantifiable measures could be used in conjunction with α and the RMS element error to determine which calibration approach is appropriate in a given scenario. The existing code's modularity allows all these ideas and many others to be easily incorporated in a controlled, logical manner.

References

- [1] Fuhrmann, D.R., "Estimation of sensor gain and phase using known field covariance," *Acoustics, Speech, and Signal Processing, 1991. ICASSP-91., 1991 International Conference on* , vol., no., pp.1369,1372 vol.2, 14-17 Apr 1991
- [2] Koerber, M.A; Fuhrmann, D.R., "Array calibration by Fourier series parameterization: stochastic maximum likelihood solution," *Statistical Signal and Array Processing, 1992. Conference Proceedings., IEEE Sixth SP Workshop on* , vol., no., pp.304,307, 7-9 Oct 1992
- [3] Robey, F.C.; Fuhrmann, D.R.; Krich, S.I, "Array Calibration Utilizing Clutter Scattering," *Statistical Signal and Array Processing., IEEE Seventh SP Workshop on* , vol., no., pp.317,320, 26-29 Jun 1994
- [4] Yu, J.; Krolik, J., "Adaptive phase-array calibration using MIMO radar clutter," *Radar Conference (RADAR), 2013 IEEE* , vol., no., pp.1,5, April 29 2013-May 3 2013
- [5] Aumann, H.M.; Fenn, A.J.; Willwerth, F.G., "Phased array antenna calibration and pattern prediction using mutual coupling measurements," *Antennas and Propagation, IEEE Transactions on* , vol.37, no.7, pp.844,850, Jul 1989
- [6] Stark, A; Johannsen, U.; Jacob, A.F., "In-situ probes for antenna array calibration," *Microwave Conference (EuMC), 2010 European* , vol., no., pp.465,468, 28-30 Sept. 2010
- [7] Lee, K.M.; Chu, R.S.; Liu, S.C., "A performance monitoring/fault isolation and correction system of a phased array antenna using transmission-line signal injection with phase toggling method," *Antennas and Propagation Society International Symposium, 1992. AP-S. 1992 Digest. Held in Conjunction with: URSI Radio Science Meeting and Nuclear EMP Meeting., IEEE* , vol., no., pp.429,432 vol.1, 18-25 June 1992
- [8] Nicolas, J.-J., "Measurements of phased array antenna fields IN SITU: A few key aspects," *Radar Conference, 2008. RADAR '08. IEEE* , vol., no., pp.1,6, 26-30 May 2008
- [9] Dongli Liu; Jun Teng; Xiaofang Shao, "A new moveable calibration method for phased array radar," *Measurement, Information and Control (MIC), 2012 International Conference on* , vol.2, no., pp.865,868, 18-20 May 2012
- [10] Shipley, C.; Woods, D., "Mutual coupling-based calibration of phased array antennas," *Phased Array Systems and Technology, 2000. Proceedings. 2000 IEEE International Conference on* , vol., no., pp.529,532, 2000

- [11] Neidman, Y.; Shavit, R.; Bronshtein, A, "Diagnostic of phased arrays with faulty elements using the mutual coupling method," *Microwaves, Antennas & Propagation, IET* , vol.3, no.2, pp.235,241, March 2009
- [12] Bekers, D.; van Dijk, R.; van Vliet, F., "Mutual-coupling based phased-array calibration: A robust and versatile approach," *Phased Array Systems & Technology, 2013 IEEE International Symposium on* , vol., no., pp.630,637, 15-18 Oct. 2013
- [13] C. J. Fulton, "Digital Array Radar Calibration and Performance Monitoring Techniques for Direct Conversion and Dual Polarization Architectures," Ph.D. dissertation, ECE, Purdue Univ., West Lafayette, IN, 2011.

Appendix A: Populating the TR Matrix

The method to generate a TR matrix at a single frequency will be illustrated here. Consider a two by two array with no leakage. If a TR matrix is to be generated, it is necessary to first know the four by four s-parameter matrix and the four by four diagonal matrix of reflections. The full four by four coupling matrix is then created according to the equation

$$\mathbf{C}_F = \mathbf{S}(\mathbf{I} - \mathbf{\Gamma S})^{-1}. \quad (\text{A.1})$$

If the first element is transmitting and the fourth element is receiving, as illustrated in Figure 39, the full coupling matrix \mathbf{C}_F can be reduced to the coupling matrix \mathbf{C} by removing the rows and columns corresponding to the unused second and third elements. Figure 40 demonstrates this operation with letters used as dummy variables.




Fig. 39: Example of a two by two array with red representing transmit, blue representing receive, and gray representing unused elements.

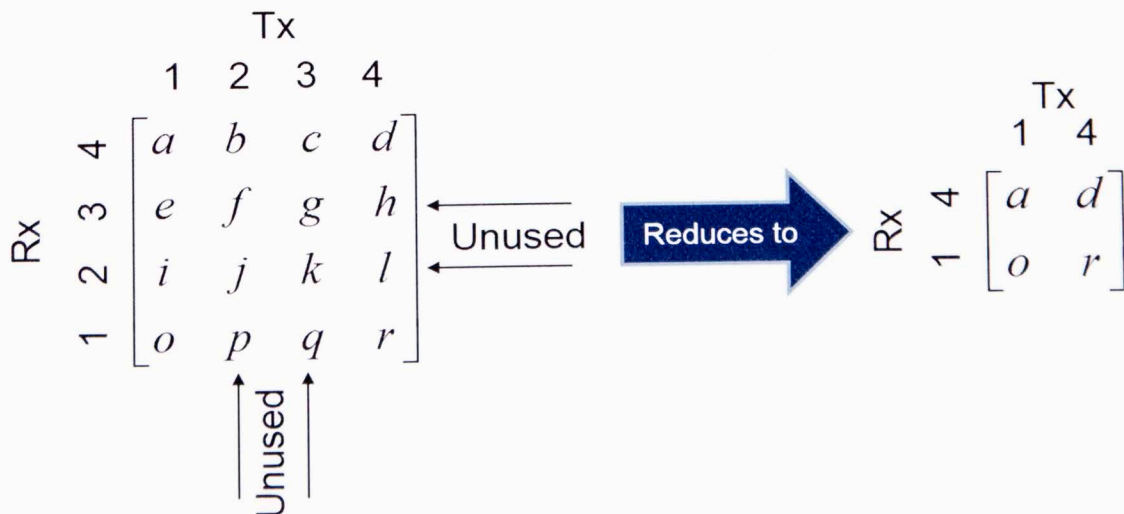


Fig. 40: Illustration of how the full coupling matrix is reduced

The diagonal matrices \mathbf{G}^R and \mathbf{G}^T are then populated with receive and transmit gains of only the used elements on the diagonals. These matrices and \mathbf{C} are then multiplied together to generate \mathbf{TR} according to the equation

$$\mathbf{TR} = \mathbf{G}^R \mathbf{C} \mathbf{G}^T. \quad (\text{A.2})$$

This final \mathbf{TR} will be an N_R by N_T matrix, where N_R is the number of active receivers and N_T is the number of active transmitters. The value at index $[r,q]$ is then the effective coupling from the r^{th} receiver to the q^{th} transmitter.

This volume is the property of the University of Oklahoma, but the literary rights of the author are a separate property and must be respected. Passages must not be copied or closely paraphrased without the previous written consent of the author. If the reader obtains any assistance from this volume, he or she must give proper credit in his own work.

I grant the University of Oklahoma Libraries permission to make a copy of my thesis/dissertation upon the request of individuals or libraries. This permission is granted with the understanding that a copy will be provided for research purposes only, and that requestors will be informed of these restrictions.

NAME  _____

DATE  _____

A library which borrows this thesis/dissertation for use by its patrons is expected to secure the signature of each user.

This thesis/dissertation by ADAM E. MITCHELL has been used by the following persons, whose signatures attest their acceptance of the above restrictions.

NAME AND ADDRESS _____ DATE