

A Unified Calibration Framework for 21 cm Cosmology

RUBY BYRNE,¹ MIGUEL F. MORALES,¹ BRYNA J. HAZELTON,^{1,2} AND MICHAEL WILENSKY¹

¹*Physics Department, University of Washington, Seattle, WA, USA*

²*eScience Institute, University of Washington, Seattle, WA, USA*

(Accepted March 2, 2021)

Submitted to Monthly Notices of the Royal Astronomical Society

ABSTRACT

Calibration precision is currently a limiting systematic in 21 cm cosmology experiments. While there are innumerable calibration approaches, most can be categorized as either ‘sky-based,’ relying on an extremely accurate model of astronomical foreground emission, or ‘redundant,’ requiring a precisely regular array with near-identical antenna response patterns. Both of these classes of calibration are inflexible to the realities of interferometric measurement. In practice, errors in the foreground model, antenna position offsets, and beam response inhomogeneities degrade calibration performance and contaminate the cosmological signal. Here we show that sky-based and redundant calibration can be unified into a highly general and physically motivated calibration framework based on a Bayesian statistical formalism. Our new framework includes sky-based and redundant calibration as special cases but can additionally support relaxing the rigid assumptions implicit in those approaches. We present simulation results demonstrating that, in a simple case, working in an intermediate regime between sky-based and redundant calibration improves calibration performance. Our framework is highly general and encompasses novel calibration approaches including techniques for calibrating compact non-redundant arrays, calibrating to incomplete sky models, and constraining calibration solutions across frequency.

Keywords: Cosmology: dark ages, reionization, first stars — Cosmology: observations — Astronomical instrumentation, methods, and techniques: methods: data analysis — Astronomical instrumentation, methods, and techniques: methods: statistical

1. INTRODUCTION

Measurement of the 21 cm cosmological power spectrum would constrain models of the Epoch of Reionization (EoR), Dark Ages, and Dark Energy. These measurements are contaminated by astrophysical foreground emission that is 4-5 orders-of-magnitude brighter than the cosmological signal. Separating the cosmological and foreground signals requires extremely precise instrumental calibration. As a result, development and characterization of precision calibration techniques for 21 cm cosmology has become an active area of research (Wieringa 1992; Pen et al. 2009; Liu et al. 2010; Kazemi et al. 2011; Kazemi & Yatawatta 2013; Grobler et al. 2014; Newburgh et al. 2014; Salvini & Wijnholds 2014; Zheng et al. 2014; Yatawatta 2015; Barry et al. 2016; Berger et al. 2016; Dillon & Parsons 2016; Ewall-Wice et al. 2016; Grobler et al. 2016; Patil et al. 2016; van Weeren et al. 2016; Wijnholds et al. 2016; Ollier et al.

2017; Dillon et al. 2018; Grobler et al. 2018; Joseph et al. 2018; Li et al. 2018; Tasse et al. 2018; Byrne et al. 2019; Kohn et al. 2019; Li et al. 2019; Orosz et al. 2019; Albert et al. 2020; Dillon et al. 2020; Joseph et al. 2020; Kern et al. 2020; Mertens et al. 2020; Sob et al. 2020).

In general, calibration approaches for cosmological 21 cm power spectrum measurements can be categorized as one of two types. ‘Sky-based’ calibration uses models of the sky and instrument to simulate data. Calibration then consists of fitting measurements to the simulation. In contrast, ‘redundant calibration’ relies on highly regular arrays with many redundant baselines measuring the same sky signal (Wieringa 1992; Liu et al. 2010; Dillon & Parsons 2016; Dillon et al. 2018; Grobler et al. 2018; Kern et al. 2020; Dillon et al. 2020). Calibration then fits redundant measurements to each other, aiming for consistency between the measurements. Many

diverse approaches fall into one of these two classes of calibration.

As the field of 21 cm cosmology pushes the limits of precision interferometric calibration it has become increasingly important that calibration frameworks mitigate error while capturing all instrumental systematics. Sky-based calibration approaches assume very good *a priori* models of the sky and instrument; model errors quickly degrade calibration and can preclude detection of the cosmological signal (Barry et al. 2016; Ewall-Wice et al. 2016). Redundant calibration is likewise vulnerable to sky model errors (Byrne et al. 2019) and experiences further calibration errors from instrument non-redundancies. Antenna position offsets and beam response inhomogeneities break redundant calibration’s assumption of baseline redundancy and produce calibration errors (Joseph et al. 2018; Li et al. 2018; Orosz et al. 2019). The field requires novel calibration approaches that are resilient to sky and instrument model errors, can capture array non-redundancies, and mitigate contamination of the cosmological signal.

Combined calibration approaches integrate aspects of sky-based and redundant calibration. The Murchison Widefield Array (MWA) Phase II is a hybrid array that incorporates regular hexagonal sub-arrays within an otherwise pseudo-random array (Wayth et al. 2018). Li et al. 2018, Li et al. 2019, and Zhang et al. 2020 combine redundant calibration of the sub-arrays with sky-based calibration for the remainder of the array. Sievers 2017 proposes a calibration algorithm that relaxes redundancy requirements for redundant calibration of a regular array. This paper expands on these ideas to present a fully unified calibration framework that is statistically rigorous, highly flexible, and physically-motivated.

In §2 we introduce a calibration formalism based on a Bayesian statistical approach. In §3 we re-derive simple sky-based and redundant calibration frameworks with a focus on delineating the implicit assumptions of those approaches. In §4 we present a novel calibration framework that unifies redundant and sky-based calibration, and §5 presents simulation results that explore some of the implications of that framework. Next, we discuss extensions to the framework that allow for calibration across frequencies (§6) and fully-polarized calibration (§7). Table 1 presents a summary of variables and expressions used throughout the paper.

2. STATISTICAL FORMALISM

In this section we present statistical underpinnings of the calibration problem. We use Bayes’ theorem to formulate a likelihood function with a statistical prior. Cal-

ibration consists of maximizing this likelihood function to calculate the instrumental response.

2.1. Bayesian Statement of the Calibration Problem

Calibration can be interpreted as a model fitting problem that fits tunable parameters to data. The fitting procedure asks, ‘what are the most likely calibration parameters given the data?’ If we define our visibility data as \mathbf{v} and our tunable calibration parameters as $\boldsymbol{\theta}$, we can write the probability of $\boldsymbol{\theta}$ given \mathbf{v} as $P(\boldsymbol{\theta}|\mathbf{v})$. This quantity is also called the ‘likelihood function.’ We can use the likelihood function to calculate maximum likelihood parameters $\hat{\boldsymbol{\theta}}$ that maximize $P(\boldsymbol{\theta}|\mathbf{v})$.

As in most model fitting problems, there is no single ‘right answer’ for the form of the likelihood. One must construct a model that appropriately represents the data and mitigates systematics that contaminate the measurement. As precision calibration is crucial to the success of 21cm cosmology, choosing a well-motivated form for the likelihood function is extremely important. This paper describes a general calibration framework and explores different physically motivated likelihood functions that could improve calibration performance for 21cm cosmology.

From Bayes’ theorem we can write the likelihood function as

$$P(\boldsymbol{\theta}|\mathbf{v}) = \frac{P(\mathbf{v}|\boldsymbol{\theta})P(\boldsymbol{\theta})}{P(\mathbf{v})} \quad (1)$$

where $P(\mathbf{v}|\boldsymbol{\theta})$ is the probability of the data \mathbf{v} given model parameters $\boldsymbol{\theta}$. $P(\boldsymbol{\theta})$ is the prior probability, i.e. the probability of $\boldsymbol{\theta}$ independent of the data. $P(\mathbf{v})$ is the marginal likelihood of \mathbf{v} and is constant across all models. We can therefore simplify the problem by considering a proportionality

$$P(\boldsymbol{\theta}|\mathbf{v}) \propto P(\mathbf{v}|\boldsymbol{\theta})P(\boldsymbol{\theta}) \quad (2)$$

and maximizing the right hand side.

Maximizing the likelihood function $P(\boldsymbol{\theta}|\mathbf{v})$ is equivalent to minimizing the negative log-likelihood. Calibration algorithms typically perform the latter procedure rather than explicitly maximizing $P(\boldsymbol{\theta}|\mathbf{v})$. We describe the negative log-likelihood as

$$L(\boldsymbol{\theta}) = -C_1 \log[P(\boldsymbol{\theta}|\mathbf{v})] + C_2, \quad (3)$$

where C_1 is an arbitrary positive constant and C_2 is an arbitrary constant of either sign. Neither constant affects the result $\hat{\boldsymbol{\theta}}$ achieved by minimizing $L(\boldsymbol{\theta})$, and we therefore choose C_1 and C_2 to simplify the form of $L(\boldsymbol{\theta})$. For the case that $P(\boldsymbol{\theta}|\mathbf{v})$ takes the form of a Gaussian distribution, the negative log-likelihood is equivalent to the least-squares cost function, also called the

Variable or Expression	Definition
\mathbf{v}	data (for interferometric data, these are the measured visibilities)
$\langle \mathbf{v} \rangle$	expectation value of the data
$\boldsymbol{\theta}$	tunable calibration parameters
$\hat{\boldsymbol{\theta}}$	maximum-likelihood calibration parameters
$\zeta(\boldsymbol{\theta})$	function that maps the calibration parameters to the model of the data
$P(\boldsymbol{\theta} \mathbf{v})$	likelihood function, equal to the probability that $\boldsymbol{\theta}$ parameterizes the model given the data \mathbf{v}
$P(\mathbf{v} \boldsymbol{\theta})$	probability that the model parameterized by $\boldsymbol{\theta}$ produces data \mathbf{v}
$P(\boldsymbol{\theta})$	prior probability of $\boldsymbol{\theta}$
$P(\mathbf{v})$	marginalized probability of \mathbf{v}
$L(\boldsymbol{\theta})$	negative log-likelihood; for a Gaussian likelihood function $L(\boldsymbol{\theta})$ is also called the chi-squared (or χ^2)
$\mathbf{l}(\boldsymbol{\theta})$	Fisher information matrix
\mathbf{g}	antenna gains; these are an example of calibration parameters $\boldsymbol{\theta}$
$\mathbf{G}(\mathbf{g})$	antenna gains written as a matrix; diagonal matrix with elements $\mathbf{g}_a \mathbf{g}_b^*$
\mathbf{u}	additional calibration parameters that are not \mathbf{g}
\mathbf{m}	the <i>a priori</i> estimator of \mathbf{u}
\mathbf{C}_T	thermal covariance matrix of the data \mathbf{v}
σ_T^2	elements of \mathbf{C}_T when \mathbf{C}_T is diagonal
\mathbf{C}_M	covariance matrix of the calibration parameters \mathbf{u}
σ_M^2	elements of \mathbf{C}_M when \mathbf{C}_M is diagonal
\mathbf{A}	matrix that maps calibration parameters \mathbf{u} to the data
N_{ant}	number of antennas
N_{red}	number of redundant baseline sets
\mathbf{x}	uv coordinate vector
$\mathbf{B}_j(\mathbf{x})$	uv response of baseline j , assuming a continuous uv plane
\mathbf{B}	matrix that encodes the uv responses of baselines
f	frequency
γ_a	calibration parameters that parameterize the gain of antenna a across frequency; these are an example of calibration parameters $\boldsymbol{\theta}$
\mathbf{V}_{ab}	2×2 visibility matrix representing the 4 polarizations measured from baseline $\{a, b\}$
\mathbf{J}_{ak}	2×2 Jones matrix for antenna a at sky location k
\mathbf{S}_k	2×2 sky coherency matrix at location k
\mathbf{E}_k	polarized electric field vector of radiation emanating from sky location k
η	delay, Fourier dual of frequency f

Table 1. A summary of variables and expressions used throughout the paper. Non-bold italicized variables represent scalars (e.g. f), bold italicized variables represent vectors (e.g. \mathbf{v}), and bold sans-serif variables represent matrices (e.g. \mathbf{A}). Italicized subscripts represent indices (e.g. \mathbf{v}_j is the j -th element of the vector \mathbf{v}) and non-italic subscripts are simply descriptive (e.g. \mathbf{u}_{orig}). \log represents the natural logarithm and we denote the identity matrix $\mathbb{1}$.

chi-squared and denoted χ^2 . For the sake of generality we do not assume that the likelihood is Gaussian. $L(\boldsymbol{\theta})$ can, but does not necessarily, take the form of a chi-squared.

This formulation of the likelihood is completely general and $\boldsymbol{\theta}$ can take any form. For example, traditional direction-independent calibration defines $\boldsymbol{\theta}$ to be per-antenna gains; direction-dependent calibration allows

these gains to assume different values at different positions on the sky. Other calibration approaches such as redundant calibration may introduce additional tunable model parameters that parameterize uncertainties on the sky model. As with the gains, calibration tunes these additional parameters to the values that maximize the likelihood function.

θ models the data according to some general function $\zeta(\theta)$. Our calibration model assumes that this function can reconstruct the expectation value of the data for some values of the tunable parameters:

$$\langle v \rangle = \zeta(\hat{\theta}), \quad (4)$$

where $\langle v \rangle$ is the expectation value of the data.

A specific calibration approach is defined by the following choices:

1. A choice of the tunable model parameters θ . These parameterize the instrument model and could parameterize the sky model as well.
2. A choice of the form of the visibility model as a function of the tunable parameters $\zeta(\theta)$.
3. A choice of the form of $P(v|\theta)$ that describes the noise properties of the data.
4. A choice of the prior on the tunable parameters $P(\theta)$. This could be a flat prior, such that $P(\theta)$ is a constant, or it could have a functional form where $\hat{\theta}$ favors an expectation value.

Once these choices are defined, calibration consists of fitting maximum-likelihood parameters $\hat{\theta}$. A final choice determines how these parameters are applied to the data to produce calibrated results. In some calibration frameworks only a subset of the fitted parameters get applied to the data. The rest of the fitted parameters are internal to the calibration optimization procedure. In general, calibration parameters related to the instrument response are involved in transforming uncalibrated data into calibrated data. Calibration parameters related to the sky model are internal to calibration.

All calibration frameworks described in this paper assume that the thermal noise properties of the data are Gaussian. The probability distribution $P(v|\theta)$ from item (3) above therefore takes the form of a Gaussian distribution. While this is the most common choice, other calibration approaches relax the assumption of Gaussianity (Kazemi & Yatawatta 2013; Ollier et al. 2017; Sob et al. 2020). The calibration frameworks in this paper can be easily extended to non-Gaussian distributions.

2.2. Quantifying Constraint of Calibration Solutions

A good parameterization of the model leads to highly constrained tunable parameters. The constraint on the tunable calibration parameters θ is quantified by the Fisher information.

The Fisher information matrix \mathbf{I} is given by the negative Hessian of the log-likelihood:

$$\mathbf{I}_{jk}(\theta) = -\frac{\partial^2 \log[P(\theta|v)]}{\partial \theta_j \partial \theta_k}. \quad (5)$$

From Equation 3 we can also express the Fisher information in terms of the negative log-likelihood, $L(\theta)$:

$$\mathbf{I}_{jk}(\theta) = \frac{\partial^2 L(\theta)}{\partial \theta_j \partial \theta_k}. \quad (6)$$

One can interpret the Fisher information as a measure of the curvature of the $L(\theta)$ hypersurface. At the maximum-likelihood point, where $\theta = \hat{\theta}$, $L(\theta)$ experiences a minimum. We therefore expect that all elements of the Fisher information matrix are non-negative:

$$\mathbf{I}_{jk}(\hat{\theta}) \geq 0. \quad (7)$$

Large values of this Fisher information indicate that $L(\theta)$ experiences a sharp minimum and the tunable parameters are highly constrained. Conversely, small Fisher information values mean the model is relatively agnostic to the parameter values and that the tunable parameters are not well-constrained.

If $\mathbf{I}_{jk}(\hat{\theta}) = 0$ for all k then the parameter θ_j is completely unconstrained by the model. It follows that the calibration solutions are degenerate. $L(\theta)$ does not experience a unique minimum but is rather minimized for any value of θ_j . In this case, θ_j is a degenerate parameter.

It is not always evident when calibration solutions are degenerate. Degenerate parameters can consist of linear combinations of the tunable parameters θ . We can calculate the number of degenerate parameters by taking the rank of the Fisher information matrix at the maximum-likelihood point. The calibration solutions are degenerate if $\mathbf{I}(\hat{\theta})$ is singular. The number of degenerate parameters is equal to the dimensionality of the null space of $\mathbf{I}(\hat{\theta})$; the degenerate parameters are the eigenvectors that span the null space.

Calibration degeneracies are a major challenge in 21cm cosmology. All degeneracies must be constrained to yield physical calibration solutions. Often, calibration consists of two stages where an initial calibration framework yields degenerate solutions and second step constrains those degeneracies (Liu et al. 2010; Zheng et al. 2014; Byrne et al. 2019; Kern et al. 2020).

3. TRADITIONAL DIRECTION-INDEPENDENT CALIBRATION APPROACHES

In this section we describe two simple direction-independent calibration frameworks, which we call ‘sky-

based calibration’ and ‘redundant calibration.’ We explain how these calibration frameworks connect with the statistical approach described in §2.

Both these calibration frameworks assume that the likelihood function is separable in frequency, time, and polarization. This means that each frequency channel, observation interval, and polarization mode can be calibrated independently, and we therefore omit explicit frequency, time, and polarization dependence. For a discussion of frequency calibration see §6; for a discussion of polarized calibration see §7. We assume per-time calibration throughout the paper.

3.1. Sky-Based Calibration

Traditional sky-based calibration makes the following choices:

1. It parameterizes the tunable instrument calibration parameters as a complex gain per antenna (and implicitly per polarization, frequency, and time step). $\boldsymbol{\theta} = \mathbf{g}$ where \mathbf{g} has length equal to the number of antennas.
2. It models the data as

$$\zeta_{ab} = \mathbf{g}_a \mathbf{g}_b^* \mathbf{m}_{ab} \quad (8)$$

where indices a and b index antennas. The combined indices ab indicate the index of the visibility formed by correlating signals from antennas a and b . \mathbf{m} are model visibilities developed with a sky model and instrument simulator. The sky model typically consists of a point-source catalog, sometimes in conjunction with a diffuse foreground emission map.

3. It describes $P(\mathbf{v}|\boldsymbol{\theta})$ as a Gaussian probability distribution:

$$P(\mathbf{v}|\boldsymbol{\theta}) \propto e^{-\frac{1}{2}[\mathbf{v}-\boldsymbol{\zeta}(\boldsymbol{\theta})]^\dagger \mathbf{C}_T^{-1}[\mathbf{v}-\boldsymbol{\zeta}(\boldsymbol{\theta})]}, \quad (9)$$

where the \dagger symbol denotes the conjugate transpose. Here \mathbf{C}_T is the thermal covariance matrix, given by

$$\mathbf{C}_T = \langle (\mathbf{v} - \langle \mathbf{v} \rangle)(\mathbf{v} - \langle \mathbf{v} \rangle)^\dagger \rangle. \quad (10)$$

Furthermore, it assumes the visibilities are independent such that \mathbf{C}_T is diagonal. The probability function therefore takes the form

$$P(\mathbf{v}|\boldsymbol{\theta}) \propto e^{-\frac{1}{2} \sum_j \frac{1}{\sigma_{T,j}^2} |\mathbf{v}_j - \boldsymbol{\zeta}_j(\boldsymbol{\theta})|^2} \quad (11)$$

where $\sigma_{T,j}^2$ are the diagonal elements of \mathbf{C}_T .

4. It uses a flat prior such that $P(\boldsymbol{\theta}) = P(\mathbf{g})$ is a constant.

With these choices we can write the likelihood function for traditional sky-based calibration as

$$P(\mathbf{g}|\mathbf{v}) \propto e^{-\frac{1}{2} \sum_{ab} \frac{1}{\sigma_{T,ab}^2} |\mathbf{v}_{ab} - \mathbf{g}_a \mathbf{g}_b^* \mathbf{m}_{ab}|^2}. \quad (12)$$

Maximizing this likelihood function is equivalent to minimizing the negative log-likelihood given by¹

$$L(\mathbf{g}) = \sum_{ab} \frac{1}{\sigma_{T,ab}^2} |\mathbf{v}_{ab} - \mathbf{g}_a \mathbf{g}_b^* \mathbf{m}_{ab}|^2. \quad (13)$$

The resulting calibration solutions are degenerate. They have one degenerate parameter (implicitly per frequency, time, and polarization) that corresponds to the overall phase of the gains. One can see this degeneracy by noting that the transformation $\mathbf{g} \rightarrow \mathbf{g}e^{i\phi}$ does not change the form of $L(\mathbf{g})$ for any real value ϕ . This phase can be interpreted as the absolute timing of incident radiation on the array and can be constrained by calibrating to a time-variable signal such as a pulsar (Pen et al. 2009).

All calibration frameworks described in this paper assume calibration to time-constant signals and therefore experience degeneracy in the overall phase of the gains. Typically this is constrained with a reference antenna. One then requires that $\text{Arg}[\hat{\mathbf{g}}_{\text{ref}}] = 0$ where $\hat{\mathbf{g}}_{\text{ref}}$ is the maximum-likelihood gain of the reference antenna and Arg denotes the complex phase. However, this means that systematic errors in the phase of the reference antenna affect all antennas. More sophisticated calibration approaches mitigate calibration errors in the reference antenna to better constrain the overall phase of the gains (Barry et al. 2019b; Li et al. 2019).

Traditional sky-based calibration assumes excellent knowledge of the sky and requires a highly complete sky model. It also assumes that the antenna responses are known up to a complex multiplicative factor. In other words, it assumes that the data can be fully modeled by the expression in Equation 8. This is clearly an inaccurate assumption. In reality, sky models are incomplete and inaccurate. These errors can propagate through the calibration process and degrade the calibration solutions (Grobler et al. 2014; Barry et al. 2016; Patil et al. 2016; Ewall-Wice et al. 2016; Joseph et al. 2020).

¹ Note that this quantity is actually twice the negative log-likelihood: $L(\mathbf{g}) = -2\log[P(\mathbf{g}|\mathbf{v})]$. However, per Equation 3 the coefficient is an arbitrary constant. For conciseness, throughout the paper we will refer to this quantity as the ‘negative log-likelihood.’

To account for sky model errors, we can replace the model visibilities with tunable calibration parameters \mathbf{u} . Now $\boldsymbol{\theta} = \{\mathbf{g}, \mathbf{u}\}$ and

$$\zeta_{ab} = \mathbf{g}_a \mathbf{g}_b^* \mathbf{u}_{ab}. \quad (14)$$

The likelihood function is now given by

$$P(\mathbf{g}, \mathbf{u} | \mathbf{v}) \propto e^{-\frac{1}{2} \sum_{ab} \frac{1}{\sigma_{T,ab}^2} |v_{ab} - \mathbf{g}_a \mathbf{g}_b^* \mathbf{u}_{ab}|^2} P(\mathbf{u}), \quad (15)$$

where $P(\mathbf{u})$ is the prior on the tunable visibility parameters. Traditional sky-based calibration assumes that \mathbf{u} is known exactly. Equation 15 converges to Equation 12 when

$$P(\mathbf{u}) \propto \delta(\mathbf{u} - \mathbf{m}), \quad (16)$$

where δ denotes the Dirac delta function. One expects that a more accurate calibration framework would not require so stringent a prior on \mathbf{u} . In §4 we explore the implications of relaxing this assumption.

3.2. Redundant Calibration

Traditional redundant calibration (Wieringa 1992) uses the following choices:

1. Like traditional sky-based calibration, traditional redundant calibration involves a tunable complex gain per antenna. However, it additionally assumes that the visibilities are not well-modeled and includes them as tunable sky model parameters. Redundant calibration works with highly regular arrays. In its traditional form, it assumes that baselines within a redundant set measure the same sky signal. The number of free visibility parameters is therefore equal to the number of redundant baseline sets. $\boldsymbol{\theta} = \{\mathbf{g}, \mathbf{u}\}$ where \mathbf{u} correspond to the visibilities from redundant baseline sets. The number of parameters to fit in calibration is equal to the number of antennas plus the number of redundant baseline sets.
2. It models the data as

$$\zeta_{ab} = \mathbf{g}_a \mathbf{g}_b^* \mathbf{A}_{ab,j} \mathbf{u}_j. \quad (17)$$

Here a and b index antennas and j indexes redundant baseline sets. \mathbf{A} is a matrix that maps \mathbf{u} to the full set of visibilities. Its elements are given by

$$\mathbf{A}_{ab,j} = \begin{cases} 1, & \text{if baseline } \{a, b\} \text{ belongs to set } j \\ 0, & \text{otherwise} \end{cases}. \quad (18)$$

3. Like traditional sky-based calibration, traditional redundant calibration describes $P(\mathbf{v} | \boldsymbol{\theta})$ as an independent Gaussian probability distribution as given by Equation 11.
4. It uses a flat prior on the tunable calibration parameters such that $P(\boldsymbol{\theta}) = P(\mathbf{g}, \mathbf{u})$ is a constant.

Under these assumptions, the likelihood function becomes

$$P(\mathbf{g}, \mathbf{u} | \mathbf{v}) \propto e^{-\frac{1}{2} \sum_{ab} \sum_j \frac{1}{\sigma_{T,ab}^2} |v_{ab} - \mathbf{g}_a \mathbf{g}_b^* \mathbf{A}_{ab,j} \mathbf{u}_j|^2}. \quad (19)$$

It follows that the negative log-likelihood is

$$L(\mathbf{g}, \mathbf{u}) = \sum_{ab} \sum_j \frac{1}{\sigma_{T,ab}^2} |v_{ab} - \mathbf{g}_a \mathbf{g}_b^* \mathbf{A}_{ab,j} \mathbf{u}_j|^2. \quad (20)$$

The additional tunable calibration parameters in traditional redundant calibration, as compared to traditional sky-based calibration, introduce additional degeneracies. Not only are the calibration solutions degenerate in the overall phase of the gains but they are also degenerate in the overall gain amplitude and the gradient of the gains' complex phase across the array (Liu et al. 2010). These degeneracies mean that redundant calibration must be combined with 'absolute calibration' to yield physical calibration solutions, where absolute calibration sets the values of the degenerate parameters. Minimizing Equation 20 is sometimes called 'relative calibration' to distinguish it from true redundant calibration, which must include both relative and absolute calibration steps (Zheng et al. 2014; Byrne et al. 2019; Kern et al. 2020).

The traditional redundant calibration framework assumes perfect redundancy of baselines within a redundant set. In order for different baselines within a redundant set to measure the same sky signal, antenna positions must lie on a perfect grid and each antennas' response must be identical up to a multiplicative gain. In practice, antenna position errors and response inhomogeneities degrade array redundancy (Joseph et al. 2018; Li et al. 2018; Orosz et al. 2019). In §4.5-§4.8 we propose new calibration frameworks that account for imperfect array redundancy.

Furthermore, traditional redundant calibration neglects covariance between redundant baseline sets. This is an appropriate assumption for some redundant arrays, such as the Donald C. Backer Precision Array for Probing the EoR (PAPER; Parsons et al. 2010; Ali et al. 2015) and the hexagonal sub-arrays in the MWA Phase II configuration (Wayth et al. 2018), but compact redundant arrays such as the Hydrogen EoR Array (HERA; DeBoer et al. 2017), the Hydrogen Intensity Real-Time

Analysis Experiment (HIRAX; Newburgh et al. 2016), and the Canadian Hydrogen Observatory and Radio-transient Detector (CHORD; Vanderlinde et al. 2019) can have appreciable covariance between baselines with centers at small uv separations. In §4.6-§4.8 we relax the assumption that baselines in different redundant baseline sets have zero covariance.

Redundant calibration can be an attractive alternative to sky-based calibration because it has reduced reliance on the sky model. One might think that redundant calibration can calibrate interferometric data without any prior knowledge of the sky — after all, Equation 20 doesn't require a sky model at all! However, this neglects the role of the sky model in absolute calibration. Although the promise of redundant calibration is to reduce the impact of sky model errors on calibration solutions, it is inaccurate to assume one can calibrate without *any* prior knowledge of the sky. Furthermore, it is possible that introducing a sky model into the relative calibration step of redundant calibration as a prior on the fit visibility values could improve calibration performance. In §4 we explore calibration approaches that do just that.

4. A GENERAL FRAMEWORK FOR DIRECTION-INDEPENDENT CALIBRATION

In this section we describe a novel and highly general calibration framework. This framework allows for extensions to traditional sky-based and redundant calibration that relax some of the non-physical assumptions of those approaches.

The calibration framework described in §4.1-§4.8 is inherently direction-independent. In §4.9 we discuss its extension to direction-dependent calibration. For a more in-depth discussion of direction-dependent calibration, see Kazemi et al. 2011, 2013; Patil et al. 2016; van Weeren et al. 2016; Tasse et al. 2018; Albert et al. 2020; or Mertens et al. 2020. Additionally, this calibration framework assumes per-frequency, per-time, and per-polarization calibration. §6 extends the framework to frequency-dependent calibration and §7 discusses fully polarized calibration techniques.

4.1. General Framework

In its most general form, this calibration framework makes the following assumptions:

1. It involves tunable complex gains per antenna (and implicitly per polarization, frequency, and time) \mathbf{g} . The gains are not direction-dependent, meaning that this calibration framework is inherently direction-independent. In addition, this calibration framework involves tunable sky parameters

\mathbf{u} . The specific form of \mathbf{u} depends on the class of calibration used. \mathbf{u} can be interpreted as visibilities, pixels in the uv plane, source flux densities, or something else altogether. In general, \mathbf{u} represents aspects of the sky model, possibly together with aspects of the instrument model, that are fit in calibration. The tunable calibration parameters are $\boldsymbol{\theta} = \{\mathbf{g}, \mathbf{u}\}$.

2. We model the data as

$$\zeta_{ab} = \mathbf{g}_a \mathbf{g}_b^* \sum_j \mathbf{A}_{ab,j} \mathbf{u}_j. \quad (21)$$

Here \mathbf{A} is a matrix that maps the sky parameters \mathbf{u} to model visibilities. In more compact matrix multiplication notation, we can write

$$\boldsymbol{\zeta} = \mathbf{G}(\mathbf{g}) \mathbf{A} \mathbf{u}, \quad (22)$$

where we have defined $\mathbf{G}(\mathbf{g})$ as a diagonal matrix with elements $\mathbf{g}_a \mathbf{g}_b^*$. In the case that \mathbf{u} represents visibilities, \mathbf{A} is simply the identity matrix: $\mathbf{A} = \mathbf{I}$.

3. We describe $P(\mathbf{v}|\boldsymbol{\theta})$ as an independent Gaussian probability distribution of the form of Equation 11.
4. We use a Gaussian prior on \mathbf{u} such that

$$P(\mathbf{u}) \propto e^{-\frac{1}{2}(\mathbf{u}-\mathbf{m})^\dagger \mathbf{C}_M^{-1}(\mathbf{u}-\mathbf{m})} \quad (23)$$

where $\mathbf{m} = \langle \mathbf{u} \rangle$. $\mathbf{C}_M = \text{cov}[\mathbf{u}, \mathbf{u}^\dagger]$ is an invertible matrix that encodes the covariances between elements of \mathbf{u} . Furthermore, this calibration framework assumes that \mathbf{g} and \mathbf{u} are independent and uses a flat prior on \mathbf{g} :

$$P(\boldsymbol{\theta}) = P(\mathbf{g}, \mathbf{u}) = P(\mathbf{g})P(\mathbf{u}) \propto P(\mathbf{u}). \quad (24)$$

Under this general framework, the likelihood function is given by

$$P(\mathbf{g}, \mathbf{u}|\mathbf{v}) \propto e^{-\frac{1}{2}[\mathbf{v}-\mathbf{G}(\mathbf{g})\mathbf{A}\mathbf{u}]^\dagger \mathbf{C}_T^{-1}[\mathbf{v}-\mathbf{G}(\mathbf{g})\mathbf{A}\mathbf{u}]} \times e^{-\frac{1}{2}(\mathbf{u}-\mathbf{m})^\dagger \mathbf{C}_M^{-1}(\mathbf{u}-\mathbf{m})} \quad (25)$$

and $L(\mathbf{g}, \mathbf{u})$ is given by

$$L(\mathbf{g}, \mathbf{u}) = [\mathbf{v} - \mathbf{G}(\mathbf{g})\mathbf{A}\mathbf{u}]^\dagger \mathbf{C}_T^{-1}[\mathbf{v} - \mathbf{G}(\mathbf{g})\mathbf{A}\mathbf{u}] + (\mathbf{u} - \mathbf{m})^\dagger \mathbf{C}_M^{-1}(\mathbf{u} - \mathbf{m}). \quad (26)$$

\mathbf{C}_T is diagonal because the thermal noise on the visibilities is assumed to be independent.

The calibration framework described here is highly abstracted. It represents a generalized treatment that can be applied to practical calibration problems. In §4.2-§4.8 we delineate examples of calibration approaches that emerge from this framework. These are just some of the possible avenues for future exploration.

4.2. Sky-Based Calibration with Partial Sky Model Knowledge

Traditional sky-based calibration, presented in §3.1, assumes that the model of the visibilities represents the true signal up to a multiplicative complex gain. Of course this is not true. Model visibilities are susceptible to errors from incomplete knowledge of the sky and instrument. We can make sky-based calibration more resilient to sky model errors by including tunable visibility parameters \mathbf{u} in calibration.

A dominant source of model visibility error stems from faint missing sources in the sky model catalog. We assume many sources exist with intensities below the sensitivity limit of the catalog. Each of these sources produces a fringe pattern in uv space. The signals from different sources add incoherently. In the limit of many sources randomly distributed on the sky, with flux densities randomly drawn from a power-law distribution, the central limit theorem dictates that the combined signal approaches Gaussian random noise in the uv plane (Tegmark et al. 1998). This suggests that the true visibilities are Gaussian-distributed around their model values and motivates using a Gaussian prior $P(\mathbf{u})$ given by Equation 23. This assumption breaks down if model visibility errors are instead dominated by a small number of bright sources that are either missing or mis-modeled in the sky model catalog. §4.3 presents a calibration framework that assumes correctly modeled source positions but uncertain source intensities.

If the baseline measurements sample discrete regions of uv space we further expect that the visibilities are independent. This means that \mathbf{C}_M is diagonal. Compact or redundant arrays can have significant baseline overlap in uv space and the assumption of independent visibilities does not hold. In §4.4–§4.8 we explore additional calibration frameworks that account for visibility covariance.

Under these assumptions, the prior on the tunable visibility parameters \mathbf{u} takes the form

$$P(\mathbf{u}) \propto e^{-\frac{1}{2} \sum_j \frac{1}{\sigma_{M,j}^2} |\mathbf{u}_j - \mathbf{m}_j|^2}, \quad (27)$$

where $\sigma_{M,j}^2$ are the diagonal elements of \mathbf{C}_M . We therefore get $L(\mathbf{g}, \mathbf{u})$ of the form

$$L(\mathbf{g}, \mathbf{u}) = \sum_{ab} \frac{1}{\sigma_{T,ab}^2} |\mathbf{v}_{ab} - \mathbf{g}_a \mathbf{g}_b^* \mathbf{u}_{ab}|^2 + \sum_j \frac{1}{\sigma_{M,j}^2} |\mathbf{u}_j - \mathbf{m}_j|^2 \quad (28)$$

and fit for calibration parameters $\boldsymbol{\theta} = \{\mathbf{g}, \mathbf{u}\}$. Note that this converges to traditional sky calibration (Equation 13) in the limit that $\sigma_{M,j}^2 \rightarrow 0$.

The addition of new tunable calibration parameters \mathbf{u} expands the number of tunable parameters as compared to traditional sky-based calibration (§3.1). Traditional sky-based calibration fits one complex parameter per antenna for a total of N_{ant} complex parameters, where N_{ant} refers to the number of antennas. Minimizing Equation 28 involves fitting an additional number of parameters equal to the number of independent baseline measurements. As autocorrelation visibilities are typically omitted from calibration, this corresponds to $(N_{\text{ant}}^2 - N_{\text{ant}})/2$ additional complex parameters. Calibration runtimes are highly dependent on the number of tunable parameters with the precise relationship depending on the implementation details.

We note that while this calibration approach involves many tunable parameters, overfitting is not a concern because the tunable visibility parameters \mathbf{u} are constrained by their prior $P(\mathbf{u})$. While their maximum-likelihood values may fit some degree of thermal noise, this noise does not impact the gains \mathbf{g} to a greater extent than it would in a traditional sky-based calibration approach.

One common sky-based calibration technique is to calibrate only on selected baselines where the sky model is trusted. For example, this could involve calibrating only on baselines longer than 50 wavelengths in order to eliminate the short baselines that are poorly modeled by a point source catalog. We can replicate this technique from the formalism described in Equation 28 by taking the limits

$$\sigma_{M,j}^2 \rightarrow \begin{cases} 0 & \text{if baseline } j \text{ is included in calibration} \\ \infty & \text{if baseline } j \text{ is excluded from calibration} \end{cases} \quad (29)$$

We now find that $\mathbf{u}_j \rightarrow \mathbf{m}_j$ for those baselines that are included in calibration: the fitted visibilities for those baselines are constrained to match the model visibilities. On the other hand, \mathbf{u}_j for baselines that are excluded from calibration are completely unconstrained. They will therefore take on values such that $\mathbf{g}_a \mathbf{g}_b^* \mathbf{u}_{ab} = \mathbf{v}_{ab}$ for any values of the gains \mathbf{g}_a and \mathbf{g}_b .

While this calibration framework can replicate binary baseline selection it need not completely include or exclude baselines. $\sigma_{M,j}^2$ quantifies the uncertainty on the model visibilities and can take any values. For example, it can be adjusted to be a function of baseline length to represent different levels of model confidence on different angular scales (Ewall-Wice et al. 2016). Rather than completely eliminating short baselines from calibration, one could instead selectively downweight them, gradually increasing the value of $\sigma_{M,j}^2$ on subsequently shorter and shorter baselines. $\sigma_{M,j}^2$ could also be calculated em-

pirically by measuring the agreement of the data and sky model as a function of baseline length.

4.3. Sky-Based Calibration with Uncertain Source Intensities

Some calibration approaches in the literature consider the case of a point source sky model in which source intensities are not well-constrained (Mitchell et al. 2008; Sievers 2017). In this case, it could be advantageous to redefine \mathbf{u} to be the source intensities rather than visibilities. Rather than assuming that sky model errors are dominated by many unmodeled sources with random positions, as in §4.2, this framework assumes that the source catalog is highly complete with well-modeled source positions. It could work well when the dominant sky model errors stem from errors in the modeled flux density of known sources. Under these assumptions, we can define $L(\mathbf{g}, \mathbf{u})$ to take the form

$$L(\mathbf{g}, \mathbf{u}) = [\mathbf{v} - \mathbf{G}(\mathbf{g})\mathbf{A}\mathbf{u}]^\dagger \mathbf{C}_T^{-1} [\mathbf{v} - \mathbf{G}(\mathbf{g})\mathbf{A}\mathbf{u}] + \sum_j \frac{1}{\sigma_{M,j}^2} |\mathbf{u}_j - \mathbf{m}_j|^2, \quad (30)$$

where $\mathbf{m}_j = \langle \mathbf{u}_j \rangle$ is the expected flux density of source j and $\sigma_{M,j}^2$ quantifies the uncertainty of that intensity. The matrix \mathbf{A} models the instrument response. The element \mathbf{A}_{jk} is equal to the contribution to visibility j from a source of unity intensity at the position of source k .

Precision calibration for 21cm cosmology often uses tens of thousands of calibrator sources in the model. Allowing the intensities of all sources to vary is computationally infeasible. Instead, a more realistic approach would be to take $\sigma_{M,j}^2 \rightarrow 0$, such that $\mathbf{u}_j \rightarrow \mathbf{m}_j$, for the majority of sources j . Then $\sigma_{M,j}^2 > 0$ only for a select subset of troublesome sources that one expects are inaccurately modeled.

This calibration approach is similar to direction-dependent calibration (§4.9) in that it can account for mis-modeled apparent source intensities that vary across the sky. However, we classify this approach as direction-independent because the fitted gains — the calibration parameters that are actually applied to the data — are not direction-dependent. Importantly, this calibration framework does not fit a per-antenna intensity variation across the sky.

4.4. Redundant Calibration With a Sky Model and Calibration of Hybrid (Redundant and Non-Redundant) Arrays

As explained in §3.2, traditional redundant calibration does not require a sky model or an explicit instrument model in the relative calibration step. However, this

means that it yields degenerate solutions. These degeneracies must be broken through an absolute calibration step that fits the degenerate parameters to model visibilities.

An extension to traditional redundant calibration incorporates model visibilities in $L(\mathbf{g}, \mathbf{u})$. This eliminates the need for a separate absolute calibration step and instead performs relative and absolute calibration simultaneously. It can also constrain the tunable visibility parameters \mathbf{u} to approximate a known sky model.

As in §3.2 we define \mathbf{u} to be the independent visibilities with length equal to the number of redundant baseline sets. We assume that different redundant baseline sets are independent and that, as in §4.2, the errors on the model uv plane are Gaussian distributed. Under these assumptions, we expand the traditional redundant calibration framework given by Equation 20 by including a Bayesian prior on \mathbf{u} :

$$L(\mathbf{g}, \mathbf{u}) = [\mathbf{v} - \mathbf{G}(\mathbf{g})\mathbf{A}\mathbf{u}]^\dagger \mathbf{C}_T^{-1} [\mathbf{v} - \mathbf{G}(\mathbf{g})\mathbf{A}\mathbf{u}] + \sum_j \frac{1}{\sigma_{M,j}^2} |\mathbf{u}_j - \mathbf{m}_j|^2. \quad (31)$$

\mathbf{A} maps the independent visibilities to the full set of visibilities and is given by Equation 18.

At this point it is useful to note that Equation 31 can be equivalently derived by considering the full basis of visibilities. Instead of defining \mathbf{u} as the visibilities from redundant baseline sets, we instead consider \mathbf{u}_{orig} , with length equal to the total number of visibilities. The associated matrix $\mathbf{A}_{\text{orig}} = \mathbf{1}$. The visibility covariance matrix $\mathbf{C}_{M,\text{orig}} = \text{cov}[\mathbf{u}_{\text{orig}}, \mathbf{u}_{\text{orig}}^\dagger]$. If we assume perfect redundancy, such that baselines within a redundant set measure the same signal, then $\text{cov}[\mathbf{u}_{\text{orig},j}, \mathbf{u}_{\text{orig},k}^*] = \text{var}[\mathbf{u}_{\text{orig},j}]$ when baselines j and k belong to the same redundant baseline set. If we further assume independence between visibilities from different redundant baseline sets then $\text{cov}[\mathbf{u}_{\text{orig},j}, \mathbf{u}_{\text{orig},k}^*] = 0$ when baselines j and k belong to different redundant baseline sets. We thereby derive a block-diagonal matrix $\mathbf{C}_{M,\text{orig}}$ where the number of blocks corresponds to the number of redundant baseline sets.

However, we would encounter a problem deriving $L(\mathbf{g}, \mathbf{u}_{\text{orig}})$ by plugging these quantities into Equation 26. $\mathbf{C}_{M,\text{orig}}$ is singular: it has rank equal to the number of redundant baseline sets. This belies that \mathbf{u}_{orig} is not a proper basis for calibration. It must be remapped into a new basis, which we derive by applying Singular Value Decomposition (SVD) to $\mathbf{C}_{M,\text{orig}}$ (see Appendix A). This generates a new calibration basis \mathbf{u} where $\mathbf{u}_{\text{orig}} = \mathbf{A}\mathbf{u}$ and \mathbf{A} is given by Equation 18. We thereby recover the calibration framework given in Equation 31.

The calibration framework described in this section unifies the ‘absolute’ and ‘relative’ steps of redundant calibration. Like traditional sky-based calibration (§3.1), it is degenerate only in the overall phase of the gains. Yet, like traditional redundant calibration (§3.2), it incorporates information from baseline redundancy.

This framework involves $N_{\text{ant}} + N_{\text{red}}$ tunable calibration parameters, where N_{red} refers to the number of redundant baseline sets. This is equivalent to the number of tunable parameters used in traditional redundant calibration.

The framework also has interesting implications for ‘hybrid’ arrays, i.e. arrays with some redundant and some non-redundant elements. Phase II of the MWA, for example, contains two hexagonal sub-arrays that support redundant calibration. These redundant sub-arrays are contained within a larger pseudo-random array. Thus far, calibrating a hybrid array has required separate calibration operations for the redundant and non-redundant elements (Li et al. 2018, 2019). In order to exploit the redundancy of the sub-arrays one had to relatively calibrate those antennas with traditional redundant calibration techniques, therefore using baseline measurements only from within the sub-arrays. The calibration framework described in this section allows for calibration of the redundant sub-array antennas using all baseline measurements. This incorporates information from all baselines involving that antenna and can increase the calibration signal-to-noise.

To explore the scope of this calibration approach we describe two opposing limits of the calibration model. The first limit takes $\sigma_{\text{M}}^2 \ll \sigma_{\text{T}}^2$. In this limit we assume that the model visibilities \mathbf{m} are well-known and require that $\mathbf{u} \approx \mathbf{m}$. Extending this limit such that $\sigma_{\text{M}}^2 \rightarrow 0$ enforces that $\mathbf{u} \rightarrow \mathbf{m}$. Equation 31 is thereby equal Equation 13, and we recover traditional sky-based calibration.

In the opposing limit $\sigma_{\text{M}}^2 \gg \sigma_{\text{T}}^2$, where σ_{T}^2 are elements of the diagonal matrix \mathbf{C}_{T} . The model visibilities \mathbf{m} will then have a negligible contribution to the relative calibration parameters, i.e. the calibration parameters constrained by minimizing Equation 20. Instead, the model visibilities will fit the absolute calibration parameters only. This limit corresponds to a unified redundant calibration framework where absolute and relative calibration are encompassed in a single minimization problem. We note that extending this limit such that $\sigma_{\text{M}}^2 \rightarrow \infty$ eliminates the second term in Equation 31. We thereby recover Equation 20 and traditional redundant calibration.

In practice, neither of these limits are physically-motivated. \mathbf{m} cannot perfectly model the visibilities;

on the other hand, it *can* give some information about the relative calibration solutions. The strength of this calibration approach is that it can inhabit the middle ground of *some* confidence in the model visibilities. As in §4.2, we could set σ_{M}^2 empirically by measuring discrepancies between the data and model. In §5 we present simulation results that explore the implications of working in an intermediate regime between traditional redundant and sky-based calibration.

4.5. Redundant Calibration Accounting for Unmodeled Imperfect Redundancy

As explained in §4.4, we can describe traditional redundant calibration with a block-diagonal \mathbf{C}_{M} when the basis of sky parameters \mathbf{u} corresponds to the full set of visibilities (denoted $\mathbf{C}_{\text{M},\text{orig}}$ in §4.4). Traditional redundant calibration assumes $\text{cov}[\mathbf{u}_j, \mathbf{u}_k^*] = \text{var}[\mathbf{u}_j]$ for visibilities \mathbf{u}_j and \mathbf{u}_k if baselines j and k belong to the same redundant set.

However, it is reasonable to assume that small antenna position errors and beam response inhomogeneities make $\text{cov}[\mathbf{u}_j, \mathbf{u}_k^*] < \text{var}[\mathbf{u}_j]$ when $j \neq k$. We can account for this by applying a suppression term to the off-diagonal elements of \mathbf{C}_{M} . This suppression factor could be calculated empirically by measuring the covariance between visibilities from redundant baselines as the array measures different fields on the sky. It could also be baseline- or antenna-dependent. For example, if one antenna is known to have a particularly irregular beam response, the baseline covariance terms for baselines that include that antenna could be preferentially suppressed.

In general, suppressing the off-diagonal elements of a block-diagonal \mathbf{C}_{M} renders it invertible (in theory at least; in practice inverting \mathbf{C}_{M} could be computationally prohibitive). We may no longer need to recast \mathbf{u} to a reduced basis such as the set of independent visibilities or the Singular Value Decomposition (SVD) basis described in §A. However, under this calibration approach the tunable visibility parameters \mathbf{u} will be highly correlated. The maximum likelihood values $\hat{\mathbf{u}}$ associated with redundant baselines will, by construction, have similar values. For this reason, we cannot use agreement of $\hat{\mathbf{u}}$ for redundant baselines as evidence of the baselines’ true degree of redundancy.

That said, considering the SVD basis can be enlightening for understanding the effect of suppressing the off-diagonal elements of \mathbf{C}_{M} by a small amount. The additional calibration degrees-of-freedom introduced correspond to differences in \mathbf{u} between visibility measurements within a redundant baseline set. The variance on these differences is small. Provided the model visibili-

ties \mathbf{m} are the same for all visibilities from a redundant set, the differences between visibility measurements in a redundant set are constrained to be near-zero.

Suppressing all off-diagonal elements of \mathbf{C}_M increases the number of tunable calibration parameters to $(N_{\text{ant}}^2 + N_{\text{ant}})/2$, as compared to $N_{\text{ant}} + N_{\text{red}}$ parameters for traditional redundant calibration. For large arrays this could substantially increase in calibration runtimes.

4.6. Redundant Calibration with Modeled Imperfect Redundancy

§4.5 describes a calibration framework that accounts for unmodeled antenna position errors and beam response inhomogeneities. However, often antenna positions and beam responses can be measured to greater accuracy than they can be controlled. For example, an antenna in a redundant array may have a known position offset from its ideal position. In this case, the covariance matrix \mathbf{C}_M can be calculated from the modeled uv responses of the antennas.

We represent the model of the uv response of baseline j as $\mathbf{B}_j(\mathbf{x})$, where \mathbf{x} is the uv position vector. These baseline response models can be developed from beam simulators or direct measurements and are readily available from instrument simulators such as Fast Holographic Deconvolution² (FHD; Sullivan et al. 2012), pyuvsim³ (Lanman et al. 2019), OSKAR⁴ (Mort et al. 2010), Precision Radio Interferometry Simulator⁵ (PRISim), or Common Astronomy Software Applications (CASA; see Jagannathan et al. 2017 for an example of using CASA with a fully-polarized primary beam model). The visibility \mathbf{u}_j relates to the true uv plane $\mathbf{S}(\mathbf{x})$ via

$$\mathbf{u}_j = \int_{-\infty}^{\infty} \mathbf{B}_j(\mathbf{x}) \mathbf{S}(\mathbf{x}) d^2 \mathbf{x}. \quad (32)$$

If we assume no uncertainty on the uv response models, the elements of the covariance matrix \mathbf{C}_M are given by

$$\begin{aligned} \mathbf{C}_{M,jk} &= \text{cov}[\mathbf{u}_j, \mathbf{u}_k^*] = \\ &= \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \mathbf{B}_j(\mathbf{x}) \mathbf{B}_k^*(\mathbf{x}') \text{cov}[\mathbf{S}(\mathbf{x}), \mathbf{S}^*(\mathbf{x}')] d^2 \mathbf{x} d^2 \mathbf{x}'. \end{aligned} \quad (33)$$

As in §4.2, we assume that sky model errors are dominated by many faint sources missing from the model catalog. We expect the resulting error in the uv plane

to be Gaussian-distributed and fully described by a variance σ_M^2 at each point. We further assume that points in the uv plane are independent:

$$\text{cov}[\mathbf{S}(\mathbf{x}), \mathbf{S}^*(\mathbf{x}')] = \begin{cases} 0 & \text{when } \mathbf{x} \neq \mathbf{x}' \\ \sigma_M^2 & \text{when } \mathbf{x} = \mathbf{x}' \end{cases}. \quad (34)$$

Under these assumptions, we get that

$$\mathbf{C}_{M,jk} = \sigma_M^2 \int_{-\infty}^{\infty} \mathbf{B}_j(\mathbf{x}) \mathbf{B}_k^*(\mathbf{x}) d^2 \mathbf{x}. \quad (35)$$

This method of calculating elements of the covariance matrix allows us to accurately represent all amounts of modeled baseline covariance in the calibration model. Redundant baselines have highly overlapping uv responses. In a closely packed array, baselines that are not constructed to be redundant may nonetheless overlap somewhat in uv coverage; this overlap introduces nonzero covariance in their measurements. Traditional redundant calibration (§3.2) assumes that $\mathbf{B}_j(\mathbf{x}) = \mathbf{B}_k(\mathbf{x})$ when baselines j and k belong to the same redundant baseline set. It also assumes that when j and k do not belong to the same redundant baseline set their uv coverage does not overlap at all: $\int_{-\infty}^{\infty} \mathbf{B}_j(\mathbf{x}) \mathbf{B}_k^*(\mathbf{x}) d^2 \mathbf{x} = 0$. If we instead calculate a covariance matrix from Equation 35 we can relax those assumptions and build a more physically-motivated calibration model that accounts for array non-redundancies and can incorporate covariances from baseline overlap in closely packed arrays such as HERA (DeBoer et al. 2017) and HIRAX (Newburgh et al. 2016).

In §4.2 we noted that a sky model may be more trusted for some baselines than others. The assumption in Equation 35 of uniform variance across the uv plane therefore must be relaxed. If we instead assume that the variance of the uv plane is approximately constant at scales equal to the size of a baseline response, we can rewrite Equation 35 as

$$\mathbf{C}_{M,jk} = \sigma_{M,j} \sigma_{M,k} \int_{-\infty}^{\infty} \mathbf{B}_j(\mathbf{x}) \mathbf{B}_k^*(\mathbf{x}) d^2 \mathbf{x}. \quad (36)$$

Here $\sigma_{M,j}^2$ is simply the variance of \mathbf{u}_j ; $\sigma_{M,j}$ is the standard deviation. As noted in §4.2, one could determine values of σ_M empirically by comparing the data to the model. Poorly-modeled power spectrum modes could be downweighted in calibration by increasing $\sigma_{M,j}$ for the associated baseline lengths. A baseline j could be excluded from calibration altogether by taking $\sigma_{M,j} \rightarrow \infty$.

We can combine the calibration approach described in this section with that of §4.5. One could expect that further non-redundancies exist beyond what is encoded in the baseline response model $\mathbf{B}_j(\mathbf{x})$. In that case, one

² <https://github.com/EoRImaging/FHD>

³ <https://github.com/RadioAstronomySoftwareGroup/pyuvsim>

⁴ <https://github.com/OxfordSKA/OSKAR>

⁵ <https://github.com/nithyanandan/PRISim>

could represent non-modeled antenna position or beam response errors by suppressing the off-diagonal elements of \mathbf{C}_M . As in §4.5, this suppression factor could be estimated from empirical measurements of the covariance of visibilities. Calculating visibility covariances across observations of different parts of the sky could help validate or adjust a covariance matrix constructed from Equation 35.

As in §4.5, under the calibration framework presented in this section the number of tunable calibration parameters lies between $N_{\text{ant}} + N_{\text{red}}$ and $(N_{\text{ant}}^2 + N_{\text{ant}})/2$, with the exact number depending on whether any baselines are modeled as fully redundant. The number of tunable parameters can be reduced by eliminating the most constrained calibration degrees-of-freedom. This is analogous to deciding that the most redundant baselines can be approximated as fully redundant. Formally, this is accomplished by performing Principal Value Decomposition (PVD) on the covariance matrix \mathbf{C}_M . We calculate the SVD basis according to Appendix A. Of the resulting parameters, we can choose not to calibrate the parameters with the smallest variances.

4.7. Covariant Calibration of Compact Non-Redundant Arrays

The calibration formalism presented in §4.6 is fully generalizable to a compact non-redundant array. The framework allows us to incorporate information from baseline covariance to constrain measurements across a non-redundant array.

For example, Phase I of the MWA (Tingay et al. 2013) has no redundant baselines. As an imaging array, it was built to be pseudo-random with maximal uv coverage, in a sense making it as non-redundant as possible (Beardsley et al. 2012). At the same time, it has a highly compact core and is uv complete out to about 50 wavelengths. This high uv coverage means that baseline measurements are not independent. Instead, many baselines are highly covariant with one another.

We can incorporate this covariance into the calibration framework by constructing \mathbf{C}_M from Equation 35. Fit visibility values \mathbf{u} would then be constrained not only by their modeled values \mathbf{m} but also by the visibilities from overlapping baselines in the uv plane.

In this way the ideas behind redundant calibration can be extended to non-redundant compact arrays. Baselines need not be highly redundant to exhibit appreciable covariance, and this covariance can be used to constrain calibration solutions. These calibration techniques have applications for many classes of interferometers, not just highly redundant arrays.

4.8. uv -Space Calibration

Traditional redundant calibration interprets the tunable sky parameters \mathbf{u} to be visibilities (denoted \mathbf{u}_{orig} in this section), however one could instead reinterpret \mathbf{u} as pixels of the uv plane. This is a natural parameter space for describing the model and means that \mathbf{C}_M does not depend on the instrument response. Instead, \mathbf{A} encodes degriding by mapping the pixels of the uv plane to the visibilities.

In §4.6 we describe the instrument response model as a continuous function in the uv plane, but in practice the uv plane is discretized. We can define a covariance matrix $\mathbf{C}_{M,\text{orig}}$ by rewriting Equation 35 for a discrete uv plane:

$$\mathbf{C}_{M,\text{orig}} = \sigma_M^2 \mathbf{B} \mathbf{B}^\dagger \quad (37)$$

where \mathbf{B} is a rectangular matrix that maps uv pixels to visibilities. \mathbf{B} has a number of columns equal to the number of pixels in the uv plane and a number of rows equal to the number of visibilities. The product $\mathbf{B} \mathbf{B}^\dagger$ is equivalent to the holographic mapping function described in Sullivan et al. 2012.

We define a new set of sky parameters \mathbf{u} that correspond to pixels in the uv plane. \mathbf{B} defines the mapping between \mathbf{u} and \mathbf{u}_{orig} :

$$\mathbf{u}_{\text{orig}} = \mathbf{B} \mathbf{u}. \quad (38)$$

We now get that

$$\begin{aligned} \mathbf{C}_{M,\text{orig}} &= \langle \mathbf{u}_{\text{orig}} \mathbf{u}_{\text{orig}}^\dagger \rangle - \langle \mathbf{u}_{\text{orig}} \rangle \langle \mathbf{u}_{\text{orig}}^\dagger \rangle \\ &= \mathbf{B} (\langle \mathbf{u} \mathbf{u}^\dagger \rangle - \langle \mathbf{u} \rangle \langle \mathbf{u}^\dagger \rangle) \mathbf{B}^\dagger \\ &= \sigma_M^2 \mathbf{B} \mathbf{B}^\dagger, \end{aligned} \quad (39)$$

so the new covariance matrix is

$$\mathbf{C}_M = \langle \mathbf{u} \mathbf{u}^\dagger \rangle - \langle \mathbf{u} \rangle \langle \mathbf{u}^\dagger \rangle = \sigma_M^2 \mathbf{1}. \quad (40)$$

The fact that this covariance matrix is diagonal highlights the assumption in §4.6 of independent uv pixels.

Plugging this new parameterization of \mathbf{u} into $L(\mathbf{g}, \mathbf{u})$ gives

$$\begin{aligned} L(\mathbf{g}, \mathbf{u}) &= [\mathbf{v} - \mathbf{G}(\mathbf{g}) \mathbf{B} \mathbf{u}]^\dagger \mathbf{C}_T^{-1} [\mathbf{v} - \mathbf{G}(\mathbf{g}) \mathbf{B} \mathbf{u}] \\ &\quad + \frac{1}{\sigma_M^2} \sum_j |\mathbf{u}_j - \mathbf{m}_j|^2, \end{aligned} \quad (41)$$

where \mathbf{m} is the model of the uv plane and j indexes over uv plane pixels. We can also relax the assumption that $\sigma_{M,j}^2 = \sigma_M^2$ for all uv pixels j and rewrite Equation 41 as

$$\begin{aligned} L(\mathbf{g}, \mathbf{u}) &= [\mathbf{v} - \mathbf{G}(\mathbf{g}) \mathbf{B} \mathbf{u}]^\dagger \mathbf{C}_T^{-1} [\mathbf{v} - \mathbf{G}(\mathbf{g}) \mathbf{B} \mathbf{u}] \\ &\quad + \sum_j \frac{1}{\sigma_{M,j}^2} |\mathbf{u}_j - \mathbf{m}_j|^2. \end{aligned} \quad (42)$$

Now σ_M^2 can represent variations in the quality of the sky model across the uv plane. A point source sky model that omits diffuse emission could have σ_M^2 increase for shorter baselines. Regions of uv space could be excluded from calibration altogether by taking the limit $\sigma_{M,j}^2 \rightarrow \infty$ for those pixels j . One could calculate σ_M^2 empirically by comparing the model of the uv plane to measurements.

We call this new calibration framework ‘ uv -space calibration.’ It offers an alternative formulation of the calibration framework described in §4.6. Parameterizing \mathbf{u} as uv pixels rather than visibilities highlights the implicit assumption of statistical independence of the pixels by making \mathbf{C}_M diagonal. Here \mathbf{m} depends on the sky model only and is calculated without an instrument simulator. The instrument model is moved from the second to the first term of $L(\mathbf{g}, \mathbf{u})$, where \mathbf{B} encodes the baseline responses to the uv plane. Baseline covariances need not be calculated explicitly. Instead, this calibration framework implicitly constrains visibilities from fully- and partially-redundant baselines from their uv plane overlap.

uv -space calibration is a natural framework for unified calibration that models all baseline covariances. As with the calibration framework described in §4.6, it can account for small non-redundancies of a redundant array due to antenna position offsets and beam inhomogeneities. Like in §4.7, it models covariances of non-redundant baselines stemming from baseline overlaps in the uv plane. Note that uv -space calibration incorporates modeled covariances only; it cannot include unmodeled imperfect redundancy as in §4.5. Finally, like the calibration framework described in §4.2, uv -space calibration can incorporate variable sky model uncertainties as a function of uv position to represent sky model incompleteness and preferentially fit calibration solutions to well-modeled uv modes.

4.9. Direction-Dependent Calibration

Direction-independent calibration assumes that the shape of the antenna responses across the sky are well-modeled *a priori*. In practice, unmodeled direction-dependent effects can degrade interferometric performance. Direction-dependent calibration is a more flexible alternative that can capture uncertainties in the spatial structure of the antenna responses.

Direction-dependent calibration has been explored extensively in the literature. In this section we show that generalized direction-dependent calibration can be formally described using the statistical framework presented in this paper. For a more thorough discussion of direction-dependent calibration see Kazemi et al. 2011,

2013; Patil et al. 2016; van Weeren et al. 2016; Tasse et al. 2018; Albert et al. 2020; and Mertens et al. 2020.

A direction-dependent extension to traditional sky-based calibration (§3.1) uses the following assumptions:

1. Like all calibration frameworks presented in this paper, direction-dependent calibration uses tunable complex gains per antenna — but now, those gains can take different values at different positions on the sky. The tunable calibration parameters are given by $\boldsymbol{\theta} = \{\mathbf{g}_1, \mathbf{g}_2, \dots, \mathbf{g}_N\}$ for N discrete sky sections, or facets. Here \mathbf{g}_n is the set of all gains for sky location n and has length equal to the number of antennas.
2. With direction-dependent calibration the visibility model $\zeta(\boldsymbol{\theta})$ no longer takes the simple form of Equation 8. As in §4, we can represent the mapping from sky parameters to visibilities with a matrix \mathbf{A} , but \mathbf{A} is now a function of the direction-dependent gains. We can write the visibility model as

$$\zeta(\mathbf{g}_1, \mathbf{g}_2, \dots, \mathbf{g}_N) = \mathbf{A}(\mathbf{g}_1, \mathbf{g}_2, \dots, \mathbf{g}_N) \mathbf{m}. \quad (43)$$

Here \mathbf{m} represents the sky model. It could be parameterized as point source intensities or the intensities of pixels on the sky. Unlike in traditional sky-based calibration, \mathbf{m} cannot be parameterized as the estimate of the visibilities because the visibilities cannot be modeled *a priori*. Simulated visibilities will depend on the direction-dependent gains.

3. As in §3.1, one can describe $P(\mathbf{v}|\boldsymbol{\theta})$ as an independent Gaussian probability distribution (Equation 11). This framework can be extended to other distributions with no loss of generality.
4. We could use a flat prior for the direction-dependent gains. Alternatively, a common prior in direction-dependent calibration requires that gains from adjacent sky sections have similar values (van Weeren et al. 2016).

Direction-dependent calibration can be a good calibration approach when antenna responses are difficult to model. The direction-dependent gains can fit response shapes that are not well-modeled and vary antenna-to-antenna. Additionally, direction-dependent calibration can be important for modeling ionospheric effects. This is particularly critical for large arrays where antennas have different lines-of-sight through the ionosphere. Refraction through the ionosphere imposes direction-dependent image distortions that vary in time as iono-

spheric conditions change (Jordan et al. 2017; Albert et al. 2020).

Direction-dependent, per-frequency calibration increases the calibration degrees-of-freedom in a way that can suppress diffuse foreground emission and the cosmological signal itself, leading to signal loss. One can protect against this signal loss by reducing calibration degrees-of-freedom across frequency (Patil et al. 2016). In §6 we explore alternatives to per-frequency calibration. With any calibration approach, one must quantify and mitigate potential signal loss stemming from coupling of the tunable calibration parameters and the cosmological signal.

5. SIMULATION RESULTS

In this section we present results from a simplified simulation exploring the novel calibration approaches described in §4. The simulation results illustrate that a unified calibration approach that combines redundant calibration with a sky model prior can improve calibration performance.

5.1. Simulation Setup

The simulation is based on an idealized redundant array of 36 antennas configured as a 6×6 square grid. The antennas are simulated as uniform circular apertures with diameters of 14 m and the array is close-packed: antenna centers are separated by 14 m. Figure 1 provides a schematic of the array configuration. We simulate visibilities from the GLEAM catalog (Hurley-Walker et al. 2017) with the pyuvsim simulation package (Lanman et al. 2019). The simulation corresponds to a zenith observation of the MWA’s ‘EoR-0’ field, a field studied extensively by the MWA’s EoR team centered at Right Ascension 0 hours and Declination -27° .

We calibrate according to the framework presented in Equation 26. We assume perfect array redundancy and take \mathbf{u} to correspond to visibilities from redundant baseline sets. \mathbf{A} therefore takes the form of Equation 18.

In contrast to traditional redundant calibration, this calibration implementation accounts for non-zero correlations between baselines from different redundant sets. As we model antennas as circular apertures, each antenna response is given by

$$\mathbf{b}(\mathbf{r}) \propto \begin{cases} 1, & \text{for } |\mathbf{r}| < a/2 \\ 0, & \text{otherwise} \end{cases}, \quad (44)$$

where \mathbf{r} is the vector position on the ground from the antenna center and a is the antenna diameter (14 m in our simulation). The baseline response is evaluated by




Figure 1. Schematic of the array layout used in simulation. The antennas are represented as uniform circular apertures with diameters of 14 m. The array comprises 36 antennas in a close-packed regular 6×6 grid.

convolving the antenna responses:

$$\mathbf{B}(\mathbf{x}) \propto \int_{-\infty}^{\infty} \mathbf{b}(\mathbf{r}) \mathbf{b}^*(\mathbf{x} - \mathbf{r}) d^2 \mathbf{r}, \quad (45)$$

where \mathbf{x} is the vector distance from the baseline center in the uv plane. Evaluating this integral and normalizing such that the integrated baseline response is equal to 1, we get that

$$\mathbf{B}(\mathbf{x}) = \frac{8}{\pi^2 a^2} \left[\cos^{-1} \left(\frac{|\mathbf{x}|}{a} \right) - \frac{|\mathbf{x}|}{a} \sqrt{1 - \left(\frac{|\mathbf{x}|}{a} \right)^2} \right] \quad (46)$$

for $|\mathbf{x}| \leq a$.

Following Equation 35, we define the baseline correlation matrix \mathbf{C}_R as

$$\mathbf{C}_{R,jk} = \text{corr}[\mathbf{u}_j, \mathbf{u}_k^*] = \frac{\int_{-\infty}^{\infty} \mathbf{B}(\mathbf{x}) \mathbf{B}(\mathbf{x} - \Delta \mathbf{x}_{jk}) d^2 \mathbf{x}}{\int_{-\infty}^{\infty} \mathbf{B}^2(\mathbf{x}) d^2 \mathbf{x}} \quad (47)$$

where $\Delta \mathbf{x}_{jk}$ is the uv separation between the centers of baselines j and k . From the definition of the correlation, the matrix is normalized such that the diagonal elements $\mathbf{C}_{R,jj} = 1$. For our array, the closest non-redundant baselines are separated by 14 m. Numerically evaluating the integrals in Equation 47, we find that those baselines have correlations of 0.1617. The next closest baseline separations are 19.80 m; those baselines have correlations of 0.0176. Subsequently more distant baselines have no uv overlap and therefore no correlation.

We define the model covariance matrix from these correlation values: $\mathbf{C}_M = \sigma_M^2 \mathbf{C}_R$, where σ_M^2 quantifies the variance on the sky model errors. The inverse quantity $\mathbf{C}_M^{-1} = \frac{1}{\sigma_M^2} \mathbf{C}_R^{-1}$. \mathbf{C}_R is full-rank; we invert it with the numpy linear algebra package.

Combining these elements, we produce $L(\mathbf{g}, \mathbf{u})$ for calibration of the form

$$L(\mathbf{g}, \mathbf{u}) = \frac{1}{\sigma_T^2} \sum_{ab} \sum_j |v_{ab} - \mathbf{g}_a \mathbf{g}_b^* \mathbf{A}_{ab,j} \mathbf{u}_j|^2 + \frac{1}{\sigma_M^2} (\mathbf{u} - \mathbf{m})^\dagger \mathbf{C}_R^{-1} (\mathbf{u} - \mathbf{m}). \quad (48)$$

As explained in §4, in the limit that $\sigma_M^2 \rightarrow \infty$ the second term of $L(\mathbf{g}, \mathbf{u})$ disappears and we recover traditional redundant calibration. In the opposite limit that $\sigma_M^2 \rightarrow 0$, $\mathbf{u} \rightarrow \mathbf{m}$ and we recover traditional sky-based calibration. We can also operate between these limits by setting σ_M^2 to a finite value that represents our confidence in the model visibilities \mathbf{m} .

Unified calibration has one degenerate parameter, corresponding to the overall phase of the gains. We constrain this term by requiring that the average complex phase of the gains is zero.

To simulate thermal noise, we inject an independent random signal into the each of the 630 simulated visibilities \mathbf{v} . The injected noise is drawn from a complex Gaussian distribution with a variance of 0.04 Jy^2 ; reflecting that, we set $\sigma_T^2 = 0.04 \text{ Jy}^2$. We explore two distinct classes of sky model errors. In §5.2 we introduce random error into the model visibilities \mathbf{m} , whereas in §5.3 we introduce errors in \mathbf{m} by omitting faint sources from the sky model. In both cases the variance on the model visibility errors is 0.16 Jy^2 , so we set $\sigma_M^2 = 0.16 \text{ Jy}^2$.

5.2. Simulating Random Model Error

In this simulation, we represent error on the model visibilities \mathbf{m} by injecting an independent random signal into each of the 60 unique baseline measurements. The error is drawn from a complex Gaussian distribution with a variance of 0.16 Jy^2 . This technique allows us to explore the implications of sky model error across a statistical ensemble of error realizations. In §5.3 we examine a more physical case in which error stems from faint missing sources in the sky model.

We implement calibration by minimizing Equation 48 with a scipy optimization routine. We perform 10,000 trials, corresponding to 100 independent realizations of each thermal noise and sky model error. Histograms of the calibrated gains $\hat{\mathbf{g}}$ over all 10,000 trials are plotted in the center column of Figure 2. The true gains $\hat{\mathbf{g}} = 1$, so deviations from 1 represent errors in the fit gains.

In addition, we calibrate the same noisy data and sky models using traditional redundant and sky-based calibration. The results from redundant calibration are plotted in the left column of Figure 2 while the results from sky-based calibration are plotted in the right column.

Of the three classes of calibration, we find that unified calibration performs best, producing the least gain error. The distribution of fit gains $\hat{\mathbf{g}}$ is more compact for unified calibration than for either redundant or sky-based calibration. Furthermore, its peak is better localized to $\hat{\mathbf{g}} = 1$. In redundant and sky-based calibration, the gain amplitudes are systematically biased low. In redundant calibration this feature results from the absolute calibration step, which constrains the overall amplitude of the gains. The bias is consistent with previous results in the field (Byrne et al. 2019) and stems from decoherence of the data and model visibilities in the presence of model errors. We explore this effect in more depth in Appendix B.

The three rows in Figure 2 depict the same underlying histograms with different overlaid contours. The contours in the top row highlight the distribution of the full set of 10,000 calibration trials. The solid white contour and the dashed white contour enclose an estimated 50% and 90% of the distribution, respectively, as calculated from a Gaussian Kernel Density Estimator (KDE).

In the second row of Figure 2 each contour describes the distribution of fit gains derived from a single realization of thermal noise. The contours enclose an estimated 50% of the distribution calculated across all 100 realizations of model error. We plot 10 such contours, randomly selected from the 100 total realizations of thermal noise. While there is some variation between the contours, the variation is small. Each realization of thermal noise produces a distribution of fit gain errors representative of the overall distribution.

In contrast, in the bottom row of Figure 2 each contour corresponds a single realization of sky model error. These contours are calculated to enclose 50% of the distribution of fit gains across all 100 realizations of thermal noise. We see that independent draws of sky model error produce very different gain distributions. This is because model error is coherent across measurements within a redundant baseline set. As regular arrays have few independent measurements of the sky signal, sky model error affecting just one baseline mode can impact calibration for a large fraction of the measured visibilities. Calibration is typically implemented for observations of a single field on the sky, so in practice we work with a single realization of sky model error. Unlike with thermal noise, the impact of sky model error on




Figure 2. 2-D histograms of the fit gains $\hat{\mathbf{g}}$ from a simulation of the array depicted in Figure 1 in the presence of random thermal noise and sky model error. The true gains are 1 (marked by a white plus sign), and deviations from 1 correspond to errors in the fit values. The three columns show the results of three different classes of calibration: traditional redundant calibration on the left, unified calibration implemented by minimizing Equation 48 in the middle, and traditional sky-based calibration on the right. The histograms depict 10,000 calibration trials corresponding to 100 separate realizations of injected thermal noise and 100 realizations of sky model error. Each row depicts identical histograms with different contours overlaid. For the top row, the white solid and dashed contours enclose an estimated 50% and 90% of the full distribution, respectively. We see that unified calibration performs better than either traditional redundant or sky-based calibration as the histogram distribution is better localized to the correct value of $\hat{\mathbf{g}} = 1$. In the middle row, each of the 10 white contours corresponds to a single realization of thermal noise. Each contour encloses 50% of the distribution of fit gains from that noise realization. In the bottom row, each of the 10 contours corresponds to a single realization of model error, again enclosing 50% of the distribution from that realization. This demonstrates that the distribution of fit gains is highly sensitive to the sky model error.

calibration solutions will not average down across long data integrations.

We can likewise compare the fit visibility values $\hat{\mathbf{u}}$ between traditional redundant calibration and unified calibration. Figure 3 presents histograms of the error in $\hat{\mathbf{u}}$. Figure 3a corresponds to traditional redundant calibration and Figure 3b corresponds to unified calibration.

Traditional sky-based calibration does not fit visibility values. Figure 3 demonstrates that unified calibration produces slightly less fit visibility error than redundant calibration, as the distribution is marginally more compact. That said, the errors in the fit gains $\hat{\mathbf{g}}$ are more relevant for evaluating calibration performance because the gains are applied to the data in calibration.




Figure 3. Complex error on fit visibility parameters $\hat{\mathbf{u}}$ from the simulation presented in Figure 2. (a) presents the results from traditional redundant calibration and (b) presents the results from unified calibration. Traditional sky-based calibration does not fit visibility parameters. The white plus signs mark the position of zero error, and the white solid and dashed contours enclose an estimated 50% and 90% of the data, respectively. Unified calibration produces slightly more precise values of $\hat{\mathbf{u}}$ than traditional redundant calibration.

These simulations illustrate that, in a simplified case, the calibration framework presented in §4 improves upon traditional calibration approaches. Traditional sky-based calibration does not fit visibility values \mathbf{u} and instead only uses the model visibilities \mathbf{m} , which are susceptible to errors from the sky model. Unified and redundant calibration approaches introduce additional calibration degrees-of-freedom through \mathbf{u} and thereby allow calibration to fit out sky model errors.

Redundant calibration uses the model visibilities \mathbf{m} for absolute calibration only. It fits the relative calibration parameters entirely from the data and imposing agreement between redundant measurements. However, thermal noise introduces error into the relative calibration parameters. Unified calibration uses the model visibilities as a prior on calibration solutions. The prior is weighted by the measured uncertainty on the sky model. This further constrains redundant calibration’s relative calibration parameters and mitigates error from thermal noise.

The calibration simulations presented in this section illustrate one implementation style of the calibration framework described in §4. The simulated instrument is perfectly redundant and we assume no errors in the instrument model. Instead, calibration errors result only from error in the model visibilities \mathbf{m} and from thermal noise. The results in Figure 2 demonstrate that, for our simplified simulation, using a unified calibration framework based on a physically-motivated instrument model produces better calibration results than traditional methods. In the next section, we demonstrate that unified calibration similarly improves calibration performance in the case that sky model errors stem from faint missing sources.

5.3. Simulating Model Incompleteness

In §5.2 we explored calibration across an ensemble of random sky model errors. As a more physically-motivated simulation, we explore the case that sky model errors stem from missing faint missing sources. This technique follows approaches used in Barry et al. 2016 and Byrne et al. 2019 and represents the fact that real source catalogs always have a completeness limit dictated by the measurement sensitivity.

To represent this source of model error we once again simulate visibilities from the GLEAM catalog, this time omitting sources with flux densities less than 100 mJy. We take these simulated visibilities to be the model visibilities \mathbf{m} . The missing faint sources introduce disparities between \mathbf{m} and the visibilities simulated with the full catalog (plotted as black ‘X’ symbols in Figure 4). The variance of these disparities is 0.16 Jy^2 , compared to the average visibility amplitude of 5.58 Jy. The black circular contour in Figure 4 has a radius equal to the standard deviation of the errors.

As in §5.3, we calibrate across 100 independent realizations of injected thermal noise drawn from a complex Gaussian distribution. The resulting calibrated gains $\hat{\mathbf{g}}$ are presented as histograms in Figure 5. Figure 5a corresponds to traditional redundant calibration, Figure 5b corresponds to unified calibration, and Figure 5c corresponds to sky-based calibration. As in Figure 2, we find that unified calibration performs better than each redundant or sky-based calibration. The fit gain distribution is most compact for unified calibration, and redundant and sky-based calibration produce gains with amplitudes that are biased low.

In Figure 6 we plot histograms of the error on the fit visibility values $\hat{\mathbf{u}}$ for redundant and unified calibration.




Figure 4. Scatter plot of the errors in the simulated model visibilities \mathbf{m} (denoted with black ‘X’ symbols). Data visibilities are simulated with the full GLEAM catalog and have an average amplitude of 5.58 Jy; model visibilities are simulated from only the GLEAM sources brighter than 100 mJy. The errors in the model visibilities stem from these missing faint sources. The white plus sign marks the position of zero error. The black solid contour has a radius equal to the standard deviation of the model visibility errors. The blue dashed contour has a radius equal to the standard deviation of the complex Gaussian distribution from which we simulate thermal noise.

As in Figure 3, we find that unified calibration produces less fit visibility error than redundant calibration.

The results of this simulation indicate that, just as unified calibration improved calibration in the presence of the random model error simulated in §5.2, it can reduce calibration error in the more physical scenario of missing faint sources in the sky model. The unified calibration framework described by Equation 48 assumes that the error on the model visibilities \mathbf{m} is Gaussian. This could be a poor approximation for some classes of sky model errors, for example, if the sky model error is dominated by mismodeling of a single bright source. Further work could develop more accurate models of sky model error based on specific knowledge of the sky. However, we expect that this assumption is appropriate for the case of sky model error from many faint missing sources distributed across the sky (and Figure 4 gives no indication that the errors on \mathbf{m} are dramatically non-Gaussian). The results in this section show that the calibration framework represented by Equation 48 can improve upon traditional calibration methods in the presence of missing sky model sources.

6. FREQUENCY CALIBRATION

In the previous sections we implicitly assume a calibration likelihood function that is separable in frequency:

$$P(\boldsymbol{\theta}|\mathbf{v}) = \prod_f P_f[\boldsymbol{\theta}(f)|\mathbf{v}(f)]. \quad (49)$$

A likelihood function of this form allows for each frequency to be calibrated separately but precludes the calibration model from incorporating cross-frequency covariance information. In this section we describe calibration extensions that break the assumption of separability of the likelihood function in frequency.

Frequency-dependent calibration errors pose a major problem for 21 cm cosmology. In particular, fast frequency structure in the calibrated gains can contaminate the cosmological signal by limiting the separability of the cosmological signal and the astronomical foregrounds.

Direction-independent gain-based calibration assumes that the instrumental response is well-modeled as a multiplicative factor on the true sky signal. In other words, a visibility is given by

$$\mathbf{v}_{ab}(f) \approx \mathbf{G}_{ab}(f)\mathbf{u}_{ab}(f) \quad (50)$$

where $\mathbf{u}_{ab}(f)$ is the true sky signal at frequency f and $\mathbf{G}_{ab}(f)$ is the instrumental gain. If calibration is error-free, one can divide the data by the gains to recover the true sky signal $\mathbf{u}_{ab}(f)$, which contains both the cosmological signal and contamination from foreground emission.

The basis of 21 cm cosmology is the principle that these signals are spectrally distinguishable. Foreground contamination is assumed to be spectrally smooth and can therefore be separated from the spectrally variant cosmological signal. If we use $\mathbf{u}_{ab}(\eta)$ to represent the Fourier transform of the signal — where η is the Fourier dual of frequency, or delay, with units time — we expect $\mathbf{u}_{ab}(\eta)$ to be free from foreground contamination at large η . The precise contaminated modes depend on the baseline length, and this contaminated region is often called the ‘wedge’ (Furlanetto et al. 2006; Datta et al. 2010; Morales & Wythe 2010; Vedantham et al. 2012; Morales et al. 2012; Parsons et al. 2012; Trott et al. 2012; Dillon et al. 2013; Hazelton et al. 2013; Thyagarajan et al. 2013, 2015).

From the convolution theorem we can rewrite Equation 50 as

$$\mathbf{v}_{ab}(\eta) \approx \mathbf{G}_{ab}(f) * \mathbf{u}_{ab}(\eta) \quad (51)$$

where $*$ denotes the convolution. Now consider the case that the gains calculated in calibration have an error that contributes to a particular delay mode η_0 . For that mode, $\mathbf{G}_{ab}(\eta_0) \rightarrow \mathbf{G}_{ab}(\eta_0) + \epsilon_{ab}(\eta_0)$, where $\epsilon_{ab}(\eta_0)$




Figure 5. 2-D histograms of the fit gains \hat{g} from a simulation of the array depicted in Figure 1. Sky model error is represented by excluding faint sources in the simulated catalog. The true gains are 1 (marked by a white plus sign), and deviations from 1 correspond to errors in the fit values. The three plots show the results of three different classes of calibration: (a) redundant calibration, (b) unified calibration, and (c) sky-based calibration. The white solid and dashed contours enclose an estimated 50% and 90% of the data, respectively. The histogram distribution in (b) is better localized to the zero error point than either (a) or (c), indicating that unified calibration delivers better calibration solutions than either traditional redundant or sky-based calibration.




Figure 6. Complex error on the fit visibility parameters \hat{u} from the simulation presented in Figure 5 for redundant calibration (a) and unified calibration (b). The white solid and dashed contours enclose an estimated 50% and 90% of the data, respectively.

is the error on the gains. Now calibration cannot recover the true sky signal; instead, it recovers the signal $\epsilon_{ab}(\eta) * u_{ab}(\eta)$. The error on the gains therefore convolves all the true sky power, including the foreground power. If η_0 is large, this means that foreground power is coupled into the modes we expect to be free from contamination, otherwise known as the ‘EoR window.’ This contaminates the cosmological signal and can preclude a detection.

For this reason, it is critical that calibration is essentially free from errors with fast frequency structure. These errors can take two forms. First of all, we require that calibration does not introduce new fast frequency structure by overfitting noise. This noise could correspond to thermal noise, which would be mitigated by extending the integration time. More importantly,

sidelobe noise from errors in the calibration sky model introduces fast frequency structure that does not average down in time. Secondly, we require that calibration accurately captures any true frequency structure in the instrumental response.

6.1. Parameterizing Gains Across Frequency

One approach to eliminating frequency-dependent calibration errors is to fit calibration solutions across frequency after-the-fact (Barry et al. 2019b; Li et al. 2019). With this approach, calibration consists of two steps. First, one calculates per-frequency calibration solutions. Next, the gains are adjusted to remove false frequency structure and fit the allowable degrees of freedom determined by a model of the instrument bandpass response.

A more optimal way of calibrating with a modeled bandpass is to fit the parameters of the bandpass di-

rectly (Yatawatta 2015; Mertens et al. 2020). We can model the gains as functions of tunable parameters γ_a for each antenna a such that

$$\mathbf{g}_a(f) = \mathbf{g}_a(f, \gamma_a). \quad (52)$$

Here γ_a could include the amplitudes and positions of known features in the bandpass or any other parameters related to fitting the instrument's bandpass response. If the bandpass is parameterized as a low-order polynomial, γ_a could be the polynomial coefficients:

$$\mathbf{g}_a(f, \gamma_a) = \sum_{n=0}^{n_{\max}} \gamma_{a,n} f^n, \quad (53)$$

where n_{\max} is the maximum mode in the low-order polynomial fit.

Because the parameters γ_a are not per-frequency, calibration can no longer be parallelized across frequency channels. Instead, the algorithm must access data from across a large frequency range at once. This increases the calibration memory requirements by a factor equal to the number of frequency channels, which can be computationally infeasible. Novel calibration algorithms allow for distributed optimization without sacrificing cross-frequency performance (Yatawatta 2015).

An extension to traditional sky-based calibration (Equation 13) that parameterizes the gains across frequencies has the following features:

1. We define tunable calibration parameters corresponding to parameterizations of the antenna gains across frequency. $\boldsymbol{\theta} = \{\gamma_1, \gamma_2, \dots\}$ where for antenna a the gain $\mathbf{g}_a(f) = \mathbf{g}_a(f, \gamma_a)$.
2. We model the data as

$$\zeta_{ab}(f) = \mathbf{g}_a(f, \gamma_a) \mathbf{g}_b^*(f, \gamma_b) \mathbf{m}_{ab}(f) \quad (54)$$

where indices a and b index antennas.

3. As in previous sections, we describe $P(\mathbf{v}|\boldsymbol{\theta})$ as an independent Gaussian probability distribution.
4. We use a flat prior such that $P(\boldsymbol{\theta}) = P(\gamma_1, \gamma_2, \dots)$ is a constant.

This version of sky-based calibration has a negative log-likelihood of the form

$$L(\gamma_1, \gamma_2, \dots) = \sum_f \sum_{ab} \frac{1}{\sigma_{T,ab}^2(f)} \times |\mathbf{v}_{ab}(f) - \mathbf{g}_a(f, \gamma_a) \mathbf{g}_b^*(f, \gamma_b) \mathbf{m}_{ab}(f)|^2. \quad (55)$$

Calibration consists of minimizing this quantity by varying the parameters γ_a for each antenna a .

This calibration approach allows for variable weighting across frequency. For example, frequency channels with greater noise contamination could be down-weighted by increasing $\sigma_T^2(f)$ with respect to the uncontaminated channels. Frequency channels can be removed from calibration altogether by taking the limit $\sigma_T^2(f) \rightarrow \infty$. If the calibration parameters are chosen to span frequencies then they will interpolate the gains across the masked frequency channels.

This approach to fitting the gains is not only applicable to sky-based calibration. One can choose to parameterize the gains across frequency in any calibration framework.

6.2. Capturing Redundancy Across Frequency

Calibrating across frequencies may allow redundant calibration to capture cross-frequency visibility covariances (Dillon & Parsons, private communication). It is possible that baselines of vastly different physical lengths are nonetheless highly covariant at different frequencies. A long baseline could measure the same sky modes at a low frequency that a short baseline measures at a higher frequency. This has interesting implications for redundant calibration. As in §6.1, this calibration approach would pose computational challenges as it precludes parallelization across frequency channels.

To illustrate cross-frequency redundant calibration we consider an extension to redundant calibration as described in §3.2. Under this extension, calibration includes the following assumptions:

1. As in §3 and §4, we parameterize the gains as per-antenna, per-frequency, and direction-independent. (Alternatively, one could combine this calibration approach with the cross-frequency gain parameterizations described in §6.1. However, for simplicity we will describe a per-frequency gain parameterization.) We parameterize the sky model with per-frequency visibilities $\mathbf{u}(f)$. The tunable calibration parameters are $\boldsymbol{\theta} = \{\mathbf{g}(f), \mathbf{u}(f)\}$.
2. As in §4, we model the data as

$$\zeta_{ab}(f) = \mathbf{g}_a(f) \mathbf{g}_b^*(f) \sum_j \mathbf{A}_{ab,j} \mathbf{u}_j(f) \quad (56)$$

for some matrix \mathbf{A} that maps \mathbf{u} to visibilities.

3. As before, $P(\mathbf{v}|\boldsymbol{\theta})$ is an independent Gaussian probability distribution (Equation 11).

4. We use a Gaussian prior on \mathbf{u} and a flat prior on \mathbf{g} such that

$$P(\boldsymbol{\theta}) \propto P(\mathbf{u}) \propto \prod_{f_1 f_2} e^{-\frac{1}{2}[\mathbf{u}(f_1) - \mathbf{m}(f_1)]^\dagger \mathbf{C}_{M, f_1 f_2}^{-1} [\mathbf{u}(f_2) - \mathbf{m}(f_2)]}. \quad (57)$$

\mathbf{C}_M now encodes covariances between every visibility at every frequency, with elements given by

$$\mathbf{C}_{M, jk f_1 f_2} = \text{cov}[\mathbf{u}_j(f_1), \mathbf{u}_k^*(f_2)] \quad (58)$$

where j and k index baselines and f_1 and f_2 index frequency channels.

$L(\mathbf{g}, \mathbf{u})$ from Equation 26 now takes the form

$$\begin{aligned} L(\mathbf{g}, \mathbf{u}) = & \sum_f \sum_{ab} \sum_j \frac{1}{\sigma_{T, ab}^2(f)} \\ & \times |v_{ab}(f) - \mathbf{g}_a(f) \mathbf{g}_b^*(f) \mathbf{A}_{ab, j} \mathbf{u}_j(f)|^2 \\ & + \sum_{f_1 f_2} \sum_{jk} [\mathbf{u}_j(f_1) - \mathbf{m}_j(f_1)]^* \mathbf{C}_{M, jk f_1 f_2}^{-1} [\mathbf{u}_k(f_2) - \mathbf{m}_k(f_2)]. \end{aligned} \quad (59)$$

This converges to traditional redundant calibration in the limit that $\mathbf{C}_{M, jk f_1 f_2} \rightarrow 0$ for $f_1 \neq f_2$.

Assuming a stable sky as a function of frequency, a physically-motivated construction of \mathbf{C}_M could follow the approach described in §4.6. From Equation 35 we get

$$\begin{aligned} \mathbf{C}_{M, jk f_1 f_2} = & \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \mathbf{B}_j(f_1, \mathbf{x}) \mathbf{B}_k^*(f_2, \mathbf{x}') \\ & \times \text{cov}[\mathbf{S}(f_1, \mathbf{x}), \mathbf{S}^*(f_2, \mathbf{x}')] d^2 \mathbf{x} d^2 \mathbf{x}', \end{aligned} \quad (60)$$

where $\mathbf{B}_j(f, \mathbf{x})$ is the uv response of baseline j and $\mathbf{S}(f, \mathbf{x})$ is the uv plane at frequency f . If, as in §4.6, we assume that different points in the uv plane are independent and that the variance across the uv plane is constant, we get that

$$\mathbf{C}_{M, jk f_1 f_2} = \mathbf{C}_{S, f_1 f_2} \int_{-\infty}^{\infty} \mathbf{B}_j(f_1, \mathbf{x}) \mathbf{B}_k^*(f_2, \mathbf{x}) d^2 \mathbf{x}. \quad (61)$$

Here $\mathbf{C}_{S, ff}$ is the variance of uv pixels at frequency f ; $\mathbf{C}_{S, f_1 f_2}$ is the covariance of uv pixels at frequencies f_1 and f_2 .

It follows from Equation 61 that the construction of a physically-motivated cross-frequency \mathbf{C}_M requires accurate modeling of the beam responses at all frequencies in question. In practice, beam responses vary considerably as a function of frequency. While baselines of different physical lengths may sample the same uv locations at different frequencies, we cannot expect their beam

responses to be the same. Therefore, redundancy is degraded across frequencies. In other words, we expect the largest elements of $\mathbf{C}_{M, f_1 f_2}$ to occur when $f_1 = f_2$.

If we assume the sky is constant across our frequency range then $\mathbf{C}_{S, f_1 f_2} = \mathbf{C}_{S, f_1 f_1}$ at all frequencies. However, in practice the sky has some frequency-dependence. We therefore expect that a more accurate model sets $\mathbf{C}_{S, f_1 f_2} < \mathbf{C}_{S, f_1 f_1}$ when $f_1 \neq f_2$. This further suppressed the elements of \mathbf{C}_M that correspond to covariances of visibilities at different frequencies.

Incorporating nonzero cross-frequency visibility covariances in calibration could be a powerful tool for constraining the frequency structure of the instrument response. However, beam shape variation across antennas and frequencies, along with frequency-dependence in the sky signal, suppress the covariance of visibilities at different frequencies. A realistic calibration model must account for these effects.

7. POLARIZED CALIBRATION

Polarized interferometry can be accomplished when each antenna has two polarization modes. This translates to four measured polarization modes from each baseline. The polarized measurement equation is given by

$$\mathbf{V}_{ab} = \sum_k \mathbf{J}_{ak} \mathbf{S}_k \mathbf{J}_{bk}^\dagger + \mathbf{N}_{ab}, \quad (62)$$

where k indexes positions on the sky. Here \mathbf{V}_{ab} is a 2×2 matrix with elements corresponding to the four polarization modes measured by baseline $\{a, b\}$. The off-diagonal elements of \mathbf{V}_{ab} correspond to the ‘cross-visibility’ formed by correlating measurements from different polarizations of the two antennas. \mathbf{J}_{ak} and \mathbf{J}_{bk} are Jones matrices associated with antennas a and b , respectively, and \mathbf{N}_{ab} gives the noise on the measurements. \mathbf{S}_k is the coherency matrix, equal to

$$\mathbf{S}_k = \langle \mathbf{E}_k \mathbf{E}_k^\dagger \rangle \quad (63)$$

where \mathbf{E}_k is the electric field vector on the sky at position k (Hamaker et al. 1996; Hamaker 2000).

\mathbf{S}_k is diagonal for an unpolarized sky, but it does not follow that \mathbf{V}_{ab} will be diagonal, even for an instrument with orthogonal antenna polarizations. Although an antenna’s two polarizations may measure orthogonal polarization modes of incident radiation from zenith (or, for a mechanically steerable antenna, from the pointing center), they will measure non-orthogonal polarization modes of off-zenith radiation. This effect is particularly pronounced for widefield instruments and fields with bright off-zenith sources. For this reason, one can expect that even unpolarized incident radiation couples into the measured cross-visibility of an interferometer.

Equation 62 is often rewritten to describe the visibilities as a vector of length 4 rather than a 2×2 matrix. In that notation, Equation 62 becomes

$$\mathbf{v}_{ab} = \sum_k (\mathbf{J}_{ak} \otimes \mathbf{J}_{bk}^*) \mathbf{s}_k + \mathbf{n}_{ab}, \quad (64)$$

where $\mathbf{J}_{ak} \otimes \mathbf{J}_{bk}^*$ is the 4×4 Kronecker product of the Jones matrices. \mathbf{v}_{ab} , \mathbf{s}_k , and \mathbf{n}_{ab} are the vector representations of the visibilities, the coherency matrix, and the noise, respectively.

7.1. Per-Polarization Calibration

Calibration frameworks presented thus far in this paper have assumed that $L(\boldsymbol{\theta})$ is separable in polarization. Under this assumption one can write the likelihood function as

$$P(\boldsymbol{\theta}|\mathbf{v}) = \prod_p P_p[\boldsymbol{\theta}_p|\mathbf{v}_{pp}], \quad (65)$$

where p indexes the instrumental polarization modes. Here $\boldsymbol{\theta}_p$ is the set of calibration parameters associated with instrumental polarization p . $\mathbf{v}_{ab,pp}$ is the (p, p) element of the matrix \mathbf{V}_{ab} ; \mathbf{v}_{pp} is the vector formed from those elements for all baselines. Note that cross-visibilitys are excluded from per-polarization calibration, which amounts to leaving potentially valuable information on the table.

We can rewrite traditional sky-based calibration (Equation 13) with explicit polarization dependence as

$$L(\mathbf{g}) = \sum_{ab} \sum_p \frac{1}{\sigma_{T,abp}^2} \left| \mathbf{v}_{ab,pp} - \mathbf{g}_{a,p} \mathbf{g}_{b,p}^* \mathbf{m}_{ab,pp} \right|^2. \quad (66)$$

Here $\mathbf{g}_{a,p}$ is the gain of the p -polarization of antenna a and $\mathbf{m}_{ab,pp}$ is an estimate of the visibility $\mathbf{v}_{ab,pp}$. The model visibilities are produced from estimations of the sky coherency matrices \mathbf{S}_k and models of the Jones matrices \mathbf{J}_{ak} .

Per-polarization calibration has a degeneracy associated with the average phase between polarizations. One can see this degeneracy by noting that the transformation $\mathbf{g}_p \rightarrow \mathbf{g}_p e^{i\phi}$ does not change the form of the $L(\mathbf{g})$, where \mathbf{g}_p is the set of all p -polarized gains. This degeneracy means that per-polarization calibration cannot calibrate the cross-visibilitys. One therefore cannot produce polarized images from per-polarization calibration. (‘Pseudo-Stokes’ I and Q images do not require the cross-visibilitys and can approximate true Stokes I and Q near zenith.)

7.2. Polarized Calibration with Per-Polarization Instrumental Gains

One approach to fully polarized calibration retains the per-polarization parameterization of the gains but incorporates the cross-visibilitys. A polarized extension to traditional sky-based calibration would use $L(\mathbf{g})$ of the form

$$L(\mathbf{g}) = \sum_{ab} \sum_{pq} \frac{1}{\sigma_{T,abpq}^2} \left| \mathbf{v}_{ab,pq} - \mathbf{g}_{a,p} \mathbf{g}_{b,q}^* \mathbf{m}_{ab,pq} \right|^2 \quad (67)$$

where p and q index instrumental polarizations.

Because this calibration framework includes the cross-visibilitys it does not have the degeneracy associated with an overall phase between polarizations. Instead, it has just one degenerate parameter corresponding to the overall phase of *all* gains. Although the gains are parameterized per-polarization, this calibration approach is fully polarized and can be used to produce polarized images.

Equation 67 represents a polarized extension to traditional sky-based calibration (Equation 13), which we have implicitly defined to be per-polarization. This polarized calibration approach could be combined with other calibration frameworks to achieve, for example, polarized redundant calibration (Kohn et al. 2019).

7.3. Calibrating Polarization Leakage

Per-polarization gain parameterization does not allow calibration to correct for polarization leakage. Polarization leakage occurs when power modeled in one visibility mode is measured as another (Sault et al. 1996). This can occur through a number of mechanisms. An antenna can be physically mis-aligned such that, for example, the so-called East-West polarized feed is somewhat skewed North-South. Antennas can experience cross-talk — for example, a signal from a North-South polarized feed could get coupled into the East-West feed’s signal path. Additionally, the ionosphere can Faraday rotate one polarization mode into another. Virtually all real antennas exhibit some level of polarization leakage. We can expand the calibration degrees-of-freedom to capture the effects.

To correct for polarization leakage in calibration, we can parameterize each of the per-antenna gains as a 2×2 matrix with two polarization indices. $\mathbf{g}_{a,pr}$ is then the gain of the p -polarization of antenna a with respect to incident r -polarized radiation. When $p \neq r$, $\mathbf{g}_{a,pr}$ encodes the polarization leakage from mode r to mode p . These off-diagonal elements of the gains matrices are often called ‘D-terms.’ An extension to traditional sky-based calibration that uses that calibration approach has

$L(\mathbf{g})$ of the form

$$L(\mathbf{g}) = \sum_{ab} \sum_{pq} \sum_{rs} \frac{1}{\sigma_{T,abpqrs}^2} \left| \mathbf{v}_{ab,pq} - \mathbf{g}_{a,pr} \mathbf{g}_{b,qs}^* \mathbf{m}_{ab,rs} \right|^2. \quad (68)$$

Here p and q index polarization modes of the measured visibilities and r and s index polarization modes of the modeled visibilities.

8. CONCLUSION

Precision calibration is essential to the success of 21 cm cosmology experiments. Low-level calibration errors contaminate the cosmological signal by limiting the separability of bright astrophysical foregrounds. As current experiments approach a detection of the 21 cm power spectrum from the EoR, novel calibration approaches must achieve new levels of precision.

While a diversity of interferometric calibration techniques exist, the current paradigm in the field delineates between sky-based and redundant calibration approaches. Sky-based calibration assumes very good *a priori* models of the sky and instrument response. It further assumes independent visibility measurements, neglecting visibility covariance from baseline response overlap in the uv plane. On the other hand, redundant calibration assumes that baselines within a redundant set are perfectly redundant both in antenna positions and response shapes, that baselines from different redundant baseline sets exhibit no covariance, and that a sky model can accurately constrain the absolute calibration parameters. The literature shows that violating these assumptions introduces calibration errors that contaminate the measurement. For example, [Barry et al. 2016](#) shows that sky model incompleteness leads to sky-based calibration errors, [Joseph et al. 2018](#) and [Orosz et al. 2019](#) explore redundant calibration errors from antenna position offsets and response irregularities, and [Byrne et al. 2019](#) examines the effect of sky model errors on the absolute calibration step of redundant calibration. Without mitigation, these errors can preclude a detection of the cosmological signal.

The standard assumptions of sky-based and redundant calibration can be inflexible to demands of calibrating real data from physical interferometers. Next-generation calibration techniques, such as those presented in [Sievers 2017](#) and this paper, combine elements of these disparate calibration approaches. New calibration frameworks must be statistically rigorous, physically-motivated, and adaptive to the complex physical systems they model. This paper presents a highly general statistical language that clarifies the assumptions implicit in a calibration formalism. It then describes new approaches to calibration that incorporate

imperfect sky models, relax redundancy requirements, and apply the ideas of redundant calibration — namely, that baselines can be covariant with one another — to a wider class of arrays than those that are typically redundantly calibrated.

While the calibration frameworks described in this paper are highly general, further extensions to this work could expand upon them. One such extension is explicitly time-dependent calibration, where the system’s evolution in time is parameterized in the calibration model. Another avenue of exploration could involve developing new approaches for constraining the overall phase of the calibration solutions. These could incorporate gated pulsar measurements, thereby augmenting sky-based or redundant calibration variants with absolute timing information.

Additional calibration extensions could constrain the frequency structure of the calibration solutions to mitigate spectral errors in calibration. [Barry et al. 2019a](#) and [Li et al. 2019](#) fit the gains’ frequency structure from the visibility autocorrelations. This protects the calibration solutions from frequency-dependent sky model errors but limits the number of visibility measurements used in calibration. As an extension to this approach, we could require that the calibration solutions fit only frequency structure that lies outside of the ‘foreground wedge.’ This approach could retain the benefits of calibrating to the autocorrelations while improving calibration signal-to-noise. We could combine this technique with regularized optimization approaches, similar to those described in [Yatawatta 2015](#), to further penalize frequency structure in the calibration solutions. This could allow calibration to fit true spectral features in the instrument’s response while mitigate spectral overfitting.

21 cm cosmology is pushing the limits of precision radio interferometry. Progress in the field will require excellent systematic suppression, including the mitigation of calibration errors. Novel calibration approaches can potentially improve measurement sensitivity. These approaches expand upon existing sky-based and redundant calibration techniques, relaxing the assumptions inherent in those approaches and incorporating more physical instrument models.

ACKNOWLEDGEMENTS

We would like to thank Henry Brinkerhoff, Ronniy Joseph, Wenyang Li, and Ian Sullivan for discussions that directly contributed to this work. We thank the anonymous reviewer who provided insightful comments that were invaluable to the development of this paper.

This work was directly supported by NSF grants AST-1613855, 1506024, 1643011, and 1835421.

DATA AVAILABILITY STATEMENT

The data underlying this article will be shared on reasonable request to the corresponding author.

REFERENCES

- Albert J. G., van Weeren R. J., Intema H. T., Röttgering H. J. A., 2020, *Astron. Astrophys.*, 635, A147
- Ali Z. S., et al., 2015, *Astrophys. J.*, 809, 61
- Barry N., Hazelton B., Sullivan I., Morales M. F., Pober J. C., 2016, *Mon. Not. R. Astron. Soc.*, 461, 3135
- Barry N., Beardsley A. P., Byrne R., Hazelton B., Morales M. F., Pober J. C., Sullivan I., 2019a, *Publ. Astron. Soc. Aust.*, 36, E026
- Barry N., et al., 2019b, *Astrophys. J.*, 884, 1
- Beardsley A. P., et al., 2012, *Mon. Not. R. Astron. Soc.*, 425, 1781
- Berger P., et al., 2016, *Ground-based Airborne Telesc. VI*, 9906, 99060D
- Byrne R., et al., 2019, *Astrophys. J.*, 875, 70
- Datta A., Bowman J. D., Carilli C. L., 2010, *Astrophys. J.*, 724, 526
- DeBoer D. R., et al., 2017, *Publ. Astron. Soc. Pacific*, 129, 045001
- Dillon J. S., Parsons A. R., 2016, *Astrophys. J.*, 826, 181
- Dillon J. S., Liu A., Tegmark M., 2013, *Phys. Rev. D*, 87, 043005
- Dillon J. S., et al., 2018, *Mon. Not. R. Astron. Soc.*, 477, 5670
- Dillon J. S., et al., 2020, arXiv e-prints
- Ewall-Wice A., Dillon J. S., Liu A., Hewitt J., 2016, *Mon. Not. R. Astron. Soc.*, 470, 1849
- Furlanetto S. R., Peng Oh S., Briggs F. H., 2006, *Phys. Rep.*, 433, 181
- Grobler T. L., Nunhokee C. D., Smirnov O. M., van Zyl A. J., de Bruyn A. G., 2014, *Mon. Not. R. Astron. Soc.*, 439, 4030
- Grobler T. L., Stewart A. J., Wijnholds S. J., Kenyon J. S., Smirnov O. M., 2016, *Mon. Not. R. Astron. Soc.*, 461, 2975
- Grobler T. L., Bernardi G., Kenyon J. S., Parsons A. R., Smirnov O. M., 2018, *Mon. Not. R. Astron. Soc.*, 476, 2410
- Hamaker J. P., 2000, *Astron. Astrophys. Suppl. Ser.*, 143, 515
- Hamaker J. P., Bregman J. D., Sault R. J., 1996, *Astron. Astrophys. Suppl. Ser.*, 117, 137
- Hazelton B. J., Morales M. F., Sullivan I. S., 2013, *Astrophys. J.*, 770, 156
- Hurley-Walker N., et al., 2017, *Mon. Not. R. Astron. Soc.*, 464, 1146
- Jagannathan P., Bhatnagar S., Rau U., Taylor A. R., 2017, *Astron. J.*, 154, 56
- Jordan C. H., et al., 2017, *Mon. Not. R. Astron. Soc.*, 471, 3974
- Joseph R. C., Trott C. M., Wayth R. B., 2018, *Astron. J.*, 156, 285
- Joseph R. C., Trott C. M., Wayth R. B., Nasirudin A., 2020, *Mon. Not. R. Astron. Soc.*, 492, 2017
- Kazemi S., Yatawatta S., 2013, *Mon. Not. R. Astron. Soc.*, 435, 597
- Kazemi S., Yatawatta S., Zaroubi S., Lampropoulos P., de Bruyn A. G., Koopmans L. V. E., Noordam J., 2011, *Mon. Not. R. Astron. Soc.*, 414, 1656
- Kazemi S., Yatawatta S., Zaroubi S., 2013, *Mon. Not. R. Astron. Soc.*, 430, 1457
- Kern N. S., et al., 2020, *Astrophys. J.*, 890, 122
- Kohn S. A., et al., 2019, *Astrophys. J.*, 882, 58
- Lanman A., Hazelton B., Jacobs D., Kolopanis M., Pober J., Aguirre J., Thyagarajan N., 2019, *J. Open Source Softw.*, 4, 1234
- Li W., et al., 2018, *Astrophys. J.*, 863, 170
- Li W., et al., 2019, *Astrophys. J.*, 887, 141
- Liu A., Tegmark M., Morrison S., Lutomirski A., Zaldarriaga M., 2010, *Mon. Not. R. Astron. Soc.*, 408, 1029
- Mertens F. G., et al., 2020, *Mon. Not. R. Astron. Soc.*, 493, 1662
- Mitchell D., Greenhill L., Wayth R., Sault R., Lonsdale C., Cappallo R., Morales M., Ord S., 2008, *IEEE J. Sel. Top. Signal Process.*, 2, 707
- Morales M. F., Wyithe J. S. B., 2010, *Annu. Rev. Astron. Astrophys.*, 48, 127
- Morales M. F., Hazelton B., Sullivan I., Beardsley A., 2012, *Astrophys. J.*, 752
- Mort B. J., Dulwich F., Salvini S., Adami K. Z., Jones M. E., 2010, in *IEEE Int. Symp.* pp 690–694, doi:10.1109/ARRAY.2010.5613289
- Newburgh L. B., et al., 2014, *Ground-based Airborne Telesc. V*, 9145, 91454V
- Newburgh L. B., et al., 2016, *Ground-based Airborne Telesc. VI*, 9906, 99065X

- Ollier V., El Korso M. N., Boyer R., Larzabal P., Pesavento M., 2017, *IEEE Trans. Signal Process.*, 65, 5649
- Orosz N., Dillon J. S., Ewall-Wice A., Parsons A. R., Thyagarajan N., 2019, *Mon. Not. R. Astron. Soc.*, 487, 537
- Parsons A. R., et al., 2010, *Astron. J.*, 139, 1468
- Parsons A. R., Pober J. C., Aguirre J. E., Carilli C. L., Jacobs D. C., Moore D. F., 2012, *Astrophys. J.*, 756
- Patil A. H., et al., 2016, *Mon. Not. R. Astron. Soc.*, 463, 4317
- Pen U. L., Chang T. C., Hirata C. M., Peterson J. B., Roy J., Gupta Y., Odegova J., Sigurdson K., 2009, *Mon. Not. R. Astron. Soc.*, 399, 181
- Salvini S., Wijnholds S. J., 2014, in 2014 31th URSI Gen. Assem. Sci. Symp. URSI GASS 2014. , doi:10.1109/URSIGASS.2014.6930038
- Sault R. J., Hamaker J. P., Bregman J. D., 1996, *Astron. Astrophys. Suppl. Ser.*, 117, 149
- Sievers J. L., 2017, arXiv e-prints
- Sob U. M., Bester H. L., Smirnov O. M., Kenyon J. S., Grobler T. L., 2020, *Mon. Not. R. Astron. Soc.*, 491, 1026
- Sullivan I. S., et al., 2012, *Astrophys. J.*, 759
- Tasse C., et al., 2018, *Astron. Astrophys.*, 611
- Tegmark M., Hamilton A. J. S., Strauss M. A., Vogeley M. S., Szalay A. S., 1998, *Astrophys. J.*, 499, 555
- Thyagarajan N., et al., 2013, *Astrophys. J.*, 776
- Thyagarajan N., et al., 2015, *Astrophys. J.*, 804
- Tingay S. J., et al., 2013, *Publ. Astron. Soc. Aust.*, 30
- Trott C. M., Wayth R. B., Tingay S. J., 2012, *Astrophys. J.*, 757
- Vanderlinde K., et al., 2019, arXiv
- Vedantham H., Udaya Shankar N., Subrahmanyam R., 2012, *Astrophys. J.*, 745
- Wayth R. B., et al., 2018, *Publ. Astron. Soc. Aust.*
- Wieringa M. H., 1992, *Exp. Astron.*, 2, 203
- Wijnholds S. J., Grobler T. L., Smirnov O. M., 2016, *Mon. Not. R. Astron. Soc.*, 457, 2331
- Yatawatta S., 2015, *Mon. Not. R. Astron. Soc.*, 449, 4506
- Zhang Z., et al., 2020, *Publications of the Astronomical Society of Australia*, 37, e045
- Zheng H., et al., 2014, *Mon. Not. R. Astron. Soc.*, 445, 1084
- van Weeren R. J., et al., 2016, *Astrophys. J. Suppl. Ser.*, 223, 2

APPENDIX

A. WORKING WITH SINGULAR COVARIANCE MATRICES

In this section we explain how to use Singular Value Decomposition (SVD) to resolve the problem of non-invertible covariance matrices by remapping the calibration parameters to a well-defined basis. This approach is fully general and can be applied to any singular covariance matrix, not just those corresponding to traditional redundant calibration as described in §3.2. It remaps an initial calibration parameter basis \mathbf{u}_{orig} , where the associated covariance matrix $\mathbf{C}_{\text{M,orig}} = \text{cov}[\mathbf{u}_{\text{orig}} \mathbf{u}_{\text{orig}}^\dagger]$ is singular, to a new basis \mathbf{u} with an invertible covariance matrix $\mathbf{C}_{\text{M}} = \text{cov}[\mathbf{u} \mathbf{u}^\dagger]$.

All covariance matrices have eigenvalues $\lambda_j \geq 0$ and eigenvectors that span the space. We can therefore transform \mathbf{u}_{orig} into the basis of eigenvectors of $\mathbf{C}_{\text{M,orig}}$. We define a new vector \mathbf{u}_{diag} as \mathbf{u}_{orig} expressed in the orthogonal eigenbasis of $\mathbf{C}_{\text{M,orig}}$:

$$\mathbf{u}_{\text{diag}} = \mathbf{E} \mathbf{u}_{\text{orig}} \quad (\text{A1})$$

where

$$\mathbf{E} = \begin{bmatrix} \mathbf{e}_1^\dagger \\ \mathbf{e}_2^\dagger \\ \vdots \\ \mathbf{e}_{N_{\text{orig}}}^\dagger \end{bmatrix} \quad (\text{A2})$$

and \mathbf{e}_j are the orthogonal eigenvectors of $\mathbf{C}_{\text{M,orig}}$. Here N_{orig} is the length of \mathbf{u}_{orig} . The associated covariance matrix is

$$\mathbf{C}_{\text{M,diag}} = \text{cov}[\mathbf{u}_{\text{diag}}, \mathbf{u}_{\text{diag}}^\dagger] = \mathbf{E} \mathbf{C}_{\text{M,orig}} \mathbf{E}^{-1} = \text{diag}(\begin{bmatrix} \lambda_1 & \lambda_2 & \cdots & \lambda_{N_{\text{orig}}} \end{bmatrix}), \quad (\text{A3})$$

where λ_j are the eigenvalues of $\mathbf{C}_{\text{M,orig}}$.

If $\mathbf{C}_{\text{M,orig}}$ is singular then it will have $\lambda_j = 0$ for some j . This has physical meaning. If $\lambda_j = 0$ then $\text{var}[\mathbf{u}_{\text{diag},j}] = 0$. It follows that those elements of \mathbf{u}_{diag} must take their expectation values: $\mathbf{u}_{\text{diag},j} = \langle \mathbf{u}_{\text{diag},j} \rangle$. We assume that $\langle \mathbf{u}_{\text{orig}} \rangle = \mathbf{m}_{\text{orig}}$, so likewise $\langle \mathbf{u}_{\text{diag}} \rangle = \mathbf{m}_{\text{diag}}$ where

$$\mathbf{m}_{\text{diag}} = \mathbf{E} \mathbf{m}_{\text{orig}}. \quad (\text{A4})$$

Therefore $\mathbf{u}_{\text{diag},j} = \mathbf{m}_{\text{diag},j}$ for all j where $\lambda_j = 0$. Those calibration parameters simply take on their modeled values; they are no longer fit during the calibration process.

We can now restrict the basis of the calibration parameters to only those that can vary in calibration. The number of independent calibration parameters is N_{red} , which is equal to the rank of $\mathbf{C}_{\text{M,orig}}$. The new basis of calibration parameters is given by

$$\mathbf{u} = \mathbf{T} \mathbf{u}_{\text{orig}}. \quad (\text{A5})$$

Here \mathbf{T} is

$$\mathbf{T} = \begin{bmatrix} \mathbf{e}_1^\dagger \\ \mathbf{e}_2^\dagger \\ \vdots \\ \mathbf{e}_{N_{\text{red}}}^\dagger \end{bmatrix}, \quad (\text{A6})$$

where we use only \mathbf{e}_j for $\lambda_j \neq 0$. $\mathbf{C}_{\text{M}} = \text{cov}[\mathbf{u}, \mathbf{u}^\dagger]$ is now diagonal and invertible.

To develop physical intuition for this remapping we can consider the case of a perfectly redundant array as described in §3.2, where \mathbf{u}_{orig} correspond to the full set of visibilities. Here $\text{cov}[\mathbf{u}_{\text{orig},j}, \mathbf{u}_{\text{orig},k}^*] = \text{var}[\mathbf{u}_{\text{orig},j}]$ if baselines j and k belong to the same redundant set and $\text{cov}[\mathbf{u}_{\text{orig},j}, \mathbf{u}_{\text{orig},k}^*] = 0$ otherwise. In this case, eigenvectors \mathbf{e}_j for $\lambda_j \neq 0$ represent the average visibilities from redundant baseline sets. Eigenvectors \mathbf{e}_j for $\lambda_j = 0$ represent visibility differences within redundant baseline sets.

\mathbf{T} is a rectangular matrix with linearly independent rows. This means that a right pseudoinverse \mathbf{T}^{p} exists such that $\mathbf{T} \mathbf{T}^{\text{p}} = \mathbb{1}$. However, it does not follow that $\mathbf{T}^{\text{p}} \mathbf{T} = \mathbb{1}$, so inverting Equation A5 is not straightforward. Rather, we get the somewhat unwieldy expression

$$\mathbf{u}_{\text{orig}} = \mathbf{T}^{\text{p}} \mathbf{u} + \sum_{\substack{j \text{ for all} \\ \lambda_j = 0}} \mathbf{m}_{\text{diag},j} \frac{\mathbf{e}_j}{|\mathbf{e}_j|^2}. \quad (\text{A7})$$

This expression simplifies when we assume that $\mathbf{m}_{\text{diag},j} = 0$ when $\lambda_j = 0$ — in other words, that the model values \mathbf{m}_{orig} are orthogonal to the null space of $\mathbf{C}_{\text{M,orig}}$. We then get that

$$\mathbf{u}_{\text{orig}} = \mathbf{T}^p \mathbf{u}. \quad (\text{A8})$$

This is a physically-motivated assumption. For the case of a perfectly redundant array, requiring that \mathbf{m}_{orig} is orthogonal to the null space means requiring that model visibilities from different baselines within redundant baseline sets are equal. Components of \mathbf{m}_{orig} in the null space indicate a disagreement between the instrument models used to produce \mathbf{m}_{orig} and $\mathbf{C}_{\text{M,orig}}$.

B. EXPLORING GAIN AMPLITUDE BIASES

In §5 we demonstrate that, in simulation, traditional sky-based and redundant calibration methods fit gain solutions $\hat{\mathbf{g}}$ that are, on average, biased low in their amplitudes. Unified calibration solutions do not exhibit this effect. In this appendix we explore the gain amplitude bias using analytic techniques. We show that suppression of the average gain amplitude in redundant and sky-based calibration stems from decoherence between the measured visibilities and the model visibilities in the presence of sky model error.

B.1. Sky-Based Calibration

Following Equation 13, traditional sky-based calibration consists of minimizing the negative log-likelihood

$$L(\mathbf{g}) = \frac{1}{\sigma_{\text{T}}^2} \sum_{\alpha} \sum_{\{b,c\} \in \alpha} |\mathbf{v}_{bc} - \mathbf{g}_b \mathbf{g}_c^* \mathbf{m}_{\alpha}|^2. \quad (\text{B9})$$

Here b and c index antennas and α indexed redundant baseline sets. $\sum_{\{b,c\} \in \alpha}$ indicates the sum over all baselines $\{b,c\}$ within redundant set α . For a non-regular array, each ‘redundant baseline set’ will contain just one baseline. We assume no baseline-dependent weighting, such that σ_{T} is a constant.

We can derive an analytic expression for the maximum likelihood solution $\hat{\mathbf{g}}$ by imposing the extremum condition

$$\left. \frac{\partial L(\mathbf{g})}{\partial \mathbf{g}_c} \right|_{\mathbf{g}=\hat{\mathbf{g}}} = 0. \quad (\text{B10})$$

Differentiating the negative log-likelihood with respect to the complex gains and evaluating gives

$$\hat{\mathbf{g}}_c = \frac{\sum_{\alpha} \sum_{b \text{ for } \{b,c\} \in \alpha} \mathbf{v}_{bc}^* \hat{\mathbf{g}}_b \mathbf{m}_{\alpha}}{\sum_{\alpha} \sum_{b \text{ for } \{b,c\} \in \alpha} |\hat{\mathbf{g}}_b \mathbf{m}_{\alpha}|^2}, \quad (\text{B11})$$

where $\sum_{b \text{ for } \{b,c\} \in \alpha}$ denotes choosing antenna b such that baseline $\{b,c\}$ belongs to redundant set α . Equation B11 represents a nonlinear system of N_{ant} coupled equations. Solving these equations together gives the calibration solutions. However, as this is not analytically tractable, practical calibration algorithms employ numerical optimization techniques.

We define a new variable A equal to the average gain amplitude

$$A = \frac{1}{N_{\text{ant}}} \sum_c |\mathbf{g}_c| \quad (\text{B12})$$

and parameterize the gains as $\mathbf{g} = A\mathbf{h}$. Plugging this into equation B11, we get

$$\hat{A}^2 = \frac{1}{N_{\text{ant}}} \sum_c \frac{\left| \sum_{\alpha} \sum_{b \text{ for } \{b,c\} \in \alpha} \mathbf{v}_{bc}^* \hat{\mathbf{h}}_b \mathbf{m}_{\alpha} \right|}{\sum_{\alpha} \sum_{b \text{ for } \{b,c\} \in \alpha} |\hat{\mathbf{h}}_b \mathbf{m}_{\alpha}|^2}. \quad (\text{B13})$$

At this point we make a series of simplifying approximations. As in the simulations in §5, we allow that the true gains are equal to 1. We can therefore make a first-order approximation of \hat{A} by setting $\hat{\mathbf{h}} = 1$:

$$\hat{A}^2 \approx \frac{1}{N_{\text{ant}}} \sum_c \frac{\left| \sum_{\alpha} \sum_{b \text{ for } \{b,c\} \in \alpha} \mathbf{v}_{bc}^* \mathbf{m}_{\alpha} \right|}{\sum_{\alpha} |\mathbf{m}_{\alpha}|^2}. \quad (\text{B14})$$

This can be rewritten as

$$\hat{A}^2 \approx \frac{|\langle \mathbf{v}^* \mathbf{m} \rangle|}{\langle |\mathbf{m}|^2 \rangle}, \quad (\text{B15})$$

where $\langle \rangle$ denotes the average across redundant baseline sets.

We describe the model visibilities \mathbf{m} as a ‘true’ visibility \mathbf{u}_{true} with error \mathbf{n}_M :

$$\mathbf{m}_\alpha = \mathbf{u}_{\text{true},\alpha} + \mathbf{n}_{M,\alpha}. \quad (\text{B16})$$

Likewise, we represent the data as

$$\mathbf{v}_{bc} = \mathbf{u}_{\text{true},\alpha} + \mathbf{n}_{T,bc} \quad (\text{B17})$$

where \mathbf{n}_T represents the thermal noise and baseline $\{b, c\}$ belongs to redundant set α . If we assume that \mathbf{n}_M , \mathbf{n}_T , and \mathbf{u}_{true} are uncorrelated with one another, then the average quantities

$$\langle \mathbf{v}^* \mathbf{m} \rangle = \langle |\mathbf{u}_{\text{true}}|^2 \rangle \quad (\text{B18})$$

and

$$\langle |\mathbf{m}|^2 \rangle = \langle |\mathbf{u}_{\text{true}}|^2 \rangle + \langle |\mathbf{n}_M|^2 \rangle. \quad (\text{B19})$$

We thereby estimate that

$$\hat{A} \approx \sqrt{\frac{\langle |\mathbf{u}_{\text{true}}|^2 \rangle}{\langle |\mathbf{u}_{\text{true}}|^2 \rangle + \langle |\mathbf{n}_M|^2 \rangle}} \leq 1. \quad (\text{B20})$$

Decoherence of the sky model error and thermal noise means that we expect suppression of average amplitude of the fit gains.

We can compare this approximation to the results of the simulation in §5.2. Across all trials in that simulation we calculate that $\langle |\mathbf{u}_{\text{true}}|^2 \rangle = 38.45 \text{ Jy}^2$ and $\langle |\mathbf{n}_M|^2 \rangle = 0.31 \text{ Jy}^2$. Plugging these into Equation B20, we estimate $A \approx 0.9960$. The average gain amplitude from the sky-based calibration simulation is 0.9964.

B.2. Redundant Calibration

The overall amplitude of the gains is degenerate in the relative calibration step of redundant calibration (Equation 20). The average gain amplitude is therefore fit in absolute calibration by minimizing the negative log-likelihood

$$L(A) = \frac{1}{\sigma_T^2} \sum_{\alpha} \sum_{\{b,c\} \in \alpha} \left| \mathbf{v}_{bc} - A^2 \hat{\mathbf{h}}_b \hat{\mathbf{h}}_c^* \mathbf{m}_\alpha \right|^2, \quad (\text{B21})$$

where $\hat{\mathbf{h}}$ have been calculated from relative calibration. For simplicity, we also assume that the gains’ phase gradient across the array has been constrained in a separate absolute calibration step. Applying the extremum condition

$$\left. \frac{\partial L(A)}{\partial A} \right|_{A=\hat{A}} = 0, \quad (\text{B22})$$

we find that

$$\hat{A}^2 = \frac{\sum_{\alpha} \sum_{\{b,c\} \in \alpha} \text{Re}(\mathbf{v}_{bc}^* \hat{\mathbf{h}}_b \hat{\mathbf{h}}_c^* \mathbf{m}_\alpha)}{\sum_{\alpha} \sum_{\{b,c\} \in \alpha} |\hat{\mathbf{h}}_b \hat{\mathbf{h}}_c^* \mathbf{m}_\alpha|^2}. \quad (\text{B23})$$

As in §B.1, we take the first-order approximation that $\hat{\mathbf{h}} = 1$, giving

$$\hat{A}^2 \approx \frac{\sum_{\alpha} \sum_{\{b,c\} \in \alpha} \text{Re}(\mathbf{v}_{bc}^* \mathbf{m}_\alpha)}{\sum_{\alpha} \sum_{\{b,c\} \in \alpha} |\mathbf{m}_\alpha|^2} \approx \frac{\langle \text{Re}(\mathbf{v}^* \mathbf{m}) \rangle}{\langle |\mathbf{m}|^2 \rangle}. \quad (\text{B24})$$

Using the decompositions in Equations B16 and B17, and once again assuming no correlation between \mathbf{n}_M , \mathbf{n}_T , and \mathbf{u}_{true} , we get a result identical to Equation B20:

$$\hat{A} \approx \sqrt{\frac{\langle |\mathbf{u}_{\text{true}}|^2 \rangle}{\langle |\mathbf{u}_{\text{true}}|^2 \rangle + \langle |\mathbf{n}_M|^2 \rangle}} \leq 1. \quad (\text{B25})$$

We expect redundant calibration to return approximately the same average gain amplitude as sky-based calibration. Indeed, in the simulation presented in §5.2, the average gain amplitude is 0.9964 for each sky-based and redundant calibration.

We note that Equation B21 represents a commonly-used absolute calibration implementation, but it is not the only method of constraining redundant calibration's gain amplitude. Future work could explore alternative approaches to absolute calibration aimed at mitigating gain amplitude bias.

B.3. Unified Calibration

From Equation 48, the unified calibration framework explored in simulation in §5 consists of minimizing the negative log-likelihood

$$L(\mathbf{g}, \mathbf{u}) = \frac{1}{\sigma_T^2} \sum_{\alpha} \sum_{\{b,c\} \in \alpha} |\mathbf{v}_{bc} - \mathbf{g}_b \mathbf{g}_c^* \mathbf{u}_{\alpha}|^2 + \frac{1}{\sigma_M^2} \sum_{\alpha} \sum_{\gamma} (\mathbf{C}_R^{-1})_{\alpha\gamma} (\mathbf{u}_{\alpha} - \mathbf{m}_{\alpha})(\mathbf{u}_{\gamma}^* - \mathbf{m}_{\gamma}^*), \quad (\text{B26})$$

where now α and γ index redundant baseline sets. For the purposes of this discussion we will approximate $\mathbf{C}_R \approx \mathbf{1}$, neglecting the covariance between redundant baseline sets. The negative log-likelihood then simplifies to

$$L(\mathbf{g}, \mathbf{u}) = \frac{1}{\sigma_T^2} \sum_{\alpha} \sum_{\{b,c\} \in \alpha} |\mathbf{v}_{bc} - \mathbf{g}_b \mathbf{g}_c^* \mathbf{u}_{\alpha}|^2 + \frac{1}{\sigma_M^2} \sum_{\alpha} |\mathbf{u}_{\alpha} - \mathbf{m}_{\alpha}|^2. \quad (\text{B27})$$

Solving the extremum condition

$$\left. \frac{\partial L(\mathbf{g}, \mathbf{u})}{\partial \mathbf{g}_c} \right|_{\mathbf{g}=\hat{\mathbf{g}}, \mathbf{u}=\hat{\mathbf{u}}} = 0. \quad (\text{B28})$$

gives

$$\hat{\mathbf{g}}_c = \frac{\sum_{\alpha} \sum_{b \text{ for } \{b,c\} \in \alpha} \mathbf{v}_{bc}^* \hat{\mathbf{g}}_b \hat{\mathbf{u}}_{\alpha}}{\sum_{\alpha} \sum_{b \text{ for } \{b,c\} \in \alpha} |\hat{\mathbf{g}}_b \hat{\mathbf{u}}_{\alpha}|^2}, \quad (\text{B29})$$

from which it follows that

$$\hat{A} = \frac{1}{N_{\text{ant}}} \sum_c \frac{\left| \sum_{\alpha} \sum_{b \text{ for } \{b,c\} \in \alpha} \mathbf{v}_{bc}^* \hat{\mathbf{g}}_b \hat{\mathbf{u}}_{\alpha} \right|}{\sum_{\alpha} \sum_{b \text{ for } \{b,c\} \in \alpha} |\hat{\mathbf{g}}_b \hat{\mathbf{u}}_{\alpha}|^2}. \quad (\text{B30})$$

We have a further extremum condition from the fit visibilities \mathbf{u} :

$$\left. \frac{\partial L(\mathbf{g}, \mathbf{u})}{\partial \mathbf{u}_{\alpha}} \right|_{\mathbf{g}=\hat{\mathbf{g}}, \mathbf{u}=\hat{\mathbf{u}}} = 0. \quad (\text{B31})$$

Solving gives

$$\hat{\mathbf{u}}_{\alpha} = \frac{\frac{1}{\sigma_T^2} \sum_{\{b,c\} \in \alpha} \mathbf{v}_{bc} \hat{\mathbf{g}}_b^* \hat{\mathbf{g}}_c + \frac{1}{\sigma_M^2} \mathbf{m}_{\alpha}}{\frac{1}{\sigma_T^2} \sum_{\{b,c\} \in \alpha} |\hat{\mathbf{g}}_b \hat{\mathbf{g}}_c^*|^2 + \frac{1}{\sigma_M^2}}. \quad (\text{B32})$$

If we again let $\mathbf{g} = A\mathbf{h}$ and take the first-order approximation $\hat{\mathbf{h}} \approx \mathbf{1}$, we get

$$\hat{A}^2 \approx \frac{1}{N_{\text{ant}}} \sum_c \frac{\left| \sum_{\alpha} \sum_{b \text{ for } \{b,c\} \in \alpha} \mathbf{v}_{bc}^* \hat{\mathbf{u}}_{\alpha} \right|}{\sum_{\alpha} |\hat{\mathbf{u}}_{\alpha}|^2} \approx \frac{|\langle \mathbf{v}^* \hat{\mathbf{u}} \rangle|}{\langle |\hat{\mathbf{u}}|^2 \rangle} \quad (\text{B33})$$

and

$$\hat{\mathbf{u}}_{\alpha} \approx \frac{\frac{1}{\sigma_T^2} \hat{A}^2 \sum_{\{b,c\} \in \alpha} \mathbf{v}_{bc} + \frac{1}{\sigma_M^2} \mathbf{m}_{\alpha}}{\frac{1}{\sigma_T^2} \hat{A}^4 N_{\alpha} + \frac{1}{\sigma_M^2}}, \quad (\text{B34})$$

where N_{α} is the number of baselines in redundant set α .

We can explore the opposing limits that $\sigma_M \rightarrow 0$ and $\sigma_M \rightarrow \infty$. For $\sigma_M \rightarrow 0$, Equation B34 gives $\hat{\mathbf{u}}_{\alpha} \approx \mathbf{m}_{\alpha}$. We then get that

$$\hat{A}^2 \approx \frac{|\langle \mathbf{v}^* \mathbf{m} \rangle|}{\langle |\mathbf{m}|^2 \rangle}, \quad (\text{B35})$$

which is identical to the solution for sky-based calibration in §B.1. In the opposite limit that $\sigma_M \rightarrow \infty$, Equation B34 becomes

$$\hat{\mathbf{u}}_\alpha \approx \frac{\sum_{\{b,c\} \in \alpha} \mathbf{v}_{bc}}{\hat{A}^2 N_\alpha}. \quad (\text{B36})$$

Following Equation B17 and again assume no correlation between \mathbf{n}_T and \mathbf{u}_{true} , we find that

$$|\langle \mathbf{v}^* \hat{\mathbf{u}} \rangle| \approx \frac{1}{\hat{A}^2} \left(\langle |\mathbf{u}_{\text{true}}|^2 \rangle + \langle |\mathbf{n}_T|^2 \rangle \right) \quad (\text{B37})$$

and

$$\langle |\hat{\mathbf{u}}|^2 \rangle \approx \frac{1}{\hat{A}^4} \left(\langle |\mathbf{u}_{\text{true}}|^2 \rangle + \langle |\mathbf{n}_T|^2 \rangle \right). \quad (\text{B38})$$

We therefore get that $\hat{A}^2 \approx \hat{A}^2$, reflecting that the overall amplitudes of the gains is degenerate in the relative calibration step of redundant calibration.

In unified calibration σ_M is finite and nonzero. In order to evaluate the average gain amplitude, we make the approximation that N_α is constant across redundant baseline sets. This is a non-physical assumption—for the simulation in §5 the number of baselines in each redundant set varies from 1 to 30—but is useful for simplifying Equation B33 and deriving a rough approximation of the average gain amplitude in unified calibration. Under this approximation, we evaluate

$$|\langle \mathbf{v}^* \hat{\mathbf{u}} \rangle| \approx \frac{\left| \frac{1}{\sigma_T^2} \hat{A}^2 \langle N \rangle \langle |\mathbf{v}|^2 \rangle + \frac{1}{\sigma_M^2} \langle \mathbf{v}^* \mathbf{m} \rangle \right|}{\frac{1}{\sigma_T^2} \hat{A}^4 \langle N \rangle + \frac{1}{\sigma_M^2}} \approx \frac{\left(\frac{1}{\sigma_T^2} \hat{A}^2 \langle N \rangle + \frac{1}{\sigma_M^2} \right) \langle |\mathbf{u}_{\text{true}}|^2 \rangle + \frac{1}{\sigma_T^2} \hat{A}^2 \langle N \rangle \langle |\mathbf{n}_T|^2 \rangle}{\frac{1}{\sigma_T^2} \hat{A}^4 \langle N \rangle + \frac{1}{\sigma_M^2}} \quad (\text{B39})$$

and

$$\begin{aligned} \langle |\hat{\mathbf{u}}|^2 \rangle &\approx \frac{\frac{1}{\sigma_T^4} \hat{A}^4 \langle N \rangle^2 \langle |\mathbf{v}|^2 \rangle + \frac{1}{\sigma_M^4} \langle |\mathbf{m}|^2 \rangle + \frac{2}{\sigma_T^2 \sigma_M^2} \hat{A}^2 \langle N \rangle \langle \text{Re}(\mathbf{v}^* \mathbf{m}) \rangle}{\left(\frac{1}{\sigma_T^2} \hat{A}^4 \langle N \rangle + \frac{1}{\sigma_M^2} \right)^2} \\ &\approx \frac{\left(\frac{1}{\sigma_T^2} \hat{A}^2 \langle N \rangle + \frac{1}{\sigma_M^2} \right)^2 \langle |\mathbf{u}_{\text{true}}|^2 \rangle + \frac{1}{\sigma_T^4} \hat{A}^4 \langle N \rangle^2 \langle |\mathbf{n}_T|^2 \rangle + \frac{1}{\sigma_M^4} \langle |\mathbf{n}_M|^2 \rangle}{\left(\frac{1}{\sigma_T^2} \hat{A}^4 \langle N \rangle + \frac{1}{\sigma_M^2} \right)^2}, \end{aligned} \quad (\text{B40})$$

where $\langle N \rangle$ is the average number of baselines in a redundant set. Combining these, we get

$$\hat{A}^2 \approx \left(\frac{1}{\sigma_T^2} \hat{A}^4 \langle N \rangle + \frac{1}{\sigma_M^2} \right) \frac{\left(\frac{1}{\sigma_T^2} \hat{A}^2 \langle N \rangle + \frac{1}{\sigma_M^2} \right) \langle |\mathbf{u}_{\text{true}}|^2 \rangle + \frac{1}{\sigma_T^2} \hat{A}^2 \langle N \rangle \langle |\mathbf{n}_T|^2 \rangle}{\left(\frac{1}{\sigma_T^2} \hat{A}^2 \langle N \rangle + \frac{1}{\sigma_M^2} \right)^2 \langle |\mathbf{u}_{\text{true}}|^2 \rangle + \frac{1}{\sigma_T^4} \hat{A}^4 \langle N \rangle^2 \langle |\mathbf{n}_T|^2 \rangle + \frac{1}{\sigma_M^4} \langle |\mathbf{n}_M|^2 \rangle}, \quad (\text{B41})$$

which can be solved numerically to evaluate \hat{A} .

As noted in §B.1, in the simulation presented in §5.2 we have $\langle |\mathbf{u}_{\text{true}}|^2 \rangle = 38.45 \text{ Jy}^2$ and $\langle |\mathbf{n}_M|^2 \rangle = 0.31 \text{ Jy}^2$. We also have $\langle |\mathbf{n}_T|^2 \rangle = 0.08 \text{ Jy}^2$, $\sigma_M = 0.4 \text{ Jy}$, $\sigma_T = 0.2 \text{ Jy}$, and $\langle N \rangle = 10.5$. Plugging in these values and numerically solving, we get $\hat{A} \approx 1.0009$. In the simulation presented in §5.2, unified calibration returns an average gain amplitude of 1.0005.