Monthly Notices of the ROYAL ASTRONOMICAL SOCIETY

MNRAS **500,** 66–81 (2021)
Advance Access publication 2020 October 22

doi:10.1093/mnras/staa3229

Calibration schemes with $O(N \log N)$ scaling for large-N radio interferometers built on a regular grid

Deepthi B. Gorthi,★† Aaron R. Parsons and Joshua S. Dillon

Department of Astronomy, University of California, Berkeley, CA 94720, USA

Accepted 2020 October 13. Received 2020 October 13; in original form 2020 May 5

ABSTRACT

Future generations of radio interferometers targeting the 21 cm signal at cosmological distances with $N \gg 1000$ antennas could face a significant computational challenge in building correlators with the traditional architecture, whose computational resource requirement scales as $\mathcal{O}(N^2)$ with array size. The fundamental output of such correlators is the cross-correlation products of all antenna pairs in the array. The FFT-correlator architecture reduces the computational resources scaling to $\mathcal{O}(N \log N)$ by computing cross-correlation products through a spatial Fourier transform. However, the output of the FFT-correlator is meaningful only when the input antenna voltages are gain- and phase-calibrated. Traditionally, interferometric calibration has used the $\mathcal{O}(N^2)$ cross-correlations produced by a standard correlator. This paper proposes two real-time calibration schemes that could work in parallel with an FFT-correlator as a self-contained $\mathcal{O}(N \log N)$ correlator system that can be scaled to large-N redundant arrays. We compare the performance and scalability of these two calibration schemes and find that they result in antenna gains whose variance decreases as $1/\log N$ with increase in the size of the array.

Key words: telescopes – instrumentation: interferometers – methods: observational – cosmology: observations.

1 INTRODUCTION

Traditional correlator architectures that have been used for most radio interferometers from the Very Large Array (VLA; Thompson et al. 1980) to the Atacama Large Millimeter Array (ALMA; Escoffier et al. 2007), require computational resources that scale as $\mathcal{O}(N^2)$ with the number of antennas. More recently, there has been a renewed interest in correlator architectures that require computational resources that scale less steeply with array size, for low-frequency radio astronomy applications that require a large collecting area. The Hydrogen Epoch of Reionization Array (HERA; DeBoer et al. 2017), the Canadian Hydrogen Intensity Mapping Experiment (CHIME; Bandura et al. 2014; Newburgh et al. 2014), the Murchison Widefield Array (MWA; Tingay et al. 2013), LOw Frequency ARray (LOFAR; van Haarlem et al. 2013), and MITEoR (Zheng et al. 2014) are all built with relatively cheap antennas that can scale to large-N arrays. At the low radio frequencies that these telescopes operate at, the signal chain can also be relatively inexpensive because cryogenic cooling of receivers is not essential. Receivers are sky-noise dominated at low radio frequencies (Ellingson et al. 2005), decreasing the need to lower thermal-noise. If the correlator architecture can also scale up to large-N arrays, it will be more cost-efficient to build the collecting area required through numerous small antennas.

In a traditional correlator architecture, the signal from every antenna is cross-correlated with the signal from every other antenna in the array. The product of individual cross-correlations are called *visibilities* and the set of visibilities from different baselines in the array is called the *visibility matrix*. Panel (a) of Fig. 1 shows the architecture of a traditional FX-correlator. The first stage performs a spectral Fourier transform, computing a spectrum of the time-varying voltage signal from antennas. The second stage computes the cross-correlation of all antenna pairs, producing a time-integrated visibility matrix. The computational resources required to generate the visibility matrix and to store and process the output data products scale as $\mathcal{O}(N^2)$ with the number of antennas in the array. For large-N arrays, this cost can dominate the entire cost of the array and has been one of the limiting factors for interferometers built in the previous decade.

For interferometers with antennas on a regular grid, Daishido et al. (1991), Tegmark & Zaldarriaga (2009), and Tegmark & Zaldarriaga (2010) have proposed FFT-correlators or FFT imagers as a potential solution to this steep scaling in cost- and computational-resources. Instead of cross-correlating antenna pairs, an FFT-correlator produces visibilities through a spatial Fourier transform. If the visibilities of redundant baselines, produced by an FX-correlator, can be averaged, these two methods are equivalent by the convolution theorem (see Tegmark & Zaldarriaga 2009, 2010). By the nature of the Fast Fourier transform algorithm (Cooley & Tukey 1965) FFT-correlators only scale as $\mathcal{O}(N \log N)$, decreasing the number of computations performed in the correlator.

An important difference between an FX-correlator and an FFT-correlator is that the latter does not preserve the full visibility matrix. The spatial Fourier transform averages redundant visibilities, which are expected to be the same, in principle, since they are visibilities measured by antenna pairs with the same displacement vector. However, in practice, redundant visibilities are different due to differences in the signal chain, structure of the dish, varying cable-lengths, etc. If antenna voltages are not calibrated, the spatial Fourier transform could result in averaging dissimilar visibilities, making post-processing correction impossible as well. Hence, antenna gain-

★ E-mail: deepthigorthi@berkeley.edu
† NSF Astronomy and Astrophysics Postdoctoral Fellow.

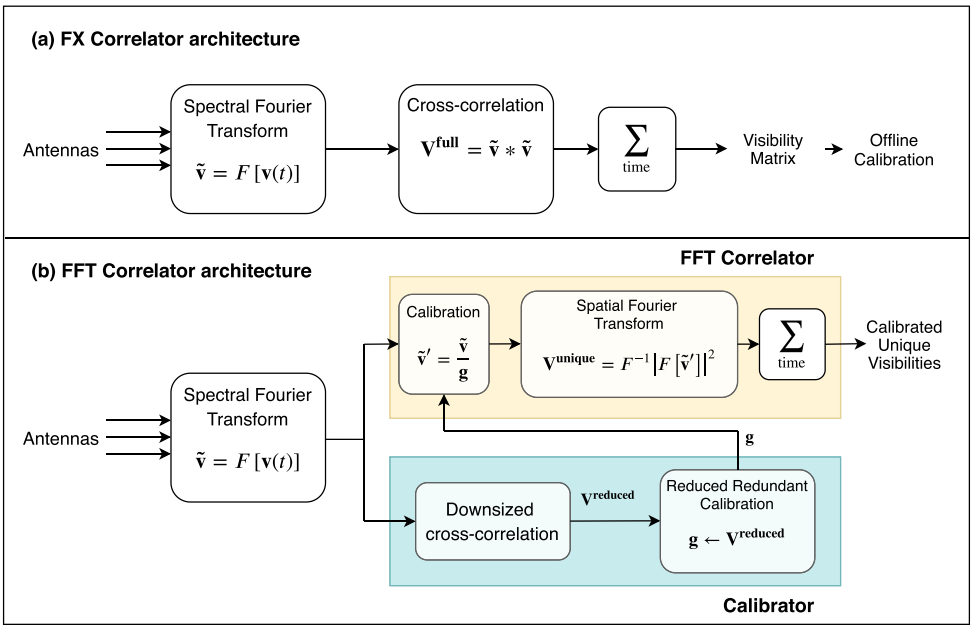

Figure 1. Panels (a) and (b) show the correlator architectures for a traditional FX-correlator and an FFT-correlator, respectively. For either architecture, the first stage Fourier transforms the voltage measured by each antenna $\boldsymbol{v}(t)$ to obtain a spectrum $\tilde{\mathbf{v}}$, and the second stage computes visibilities of all antenna pairs. For an FX-correlator, the visibility matrix $\mathbf{V}^{\mathrm{full}}$ is computed by cross-correlating the signal from every antenna with every other antenna in the array. For an FFT-correlator, the visibility matrix $\mathbf{V}^{\mathrm{unique}}$ is computed by a spatial Fourier transform on the calibrated antenna voltages (yellow box). The calibrator (blue box) operates in-parallel to the FFT-correlator and computes per-antenna gains for calibrating antenna voltages. The antenna gains are computed by performing one of the two reduced redundant-baseline calibration schemes described in this paper; on a smaller visibility matrix $\mathbf{V}^{\mathrm{reduced}}$ computed for the purpose of calibration.

and phase-calibration, prior to the spatial Fourier transform, is essential to avoid signal-loss in the FFT-correlator.

An FFT-correlator that implements the design proposed by Tegmark & Zaldarriaga (2009) is the one built by Foster et al. (2014) on the BEST-2 array at Medicina, Italy. They demonstrated that the visibilities produced by the FFT-correlator and the redundantly averaged visibilities of an FX-correlator are similar when all the antennas are calibrated before the spatial Fourier transform. However, they used a traditional FX-correlator working in parallel to generate all the visibilities required for point-source calibration. This is not a scalable solution for calibrating large-N arrays, since building an FX-correlator may not be viable.

A more generic alternative to the FFT-correlator, discussed by Thyagarajan et al. (2017), is a direct-imaging-correlator called the E-field Parallel Imaging Correlator (EPIC) that has now been deployed on the LWA (Kent et al. 2019). It works like a Modular Optimal Frequency Fourier (MOFF; Morales 2011) correlator, where antenna voltages are gridded before a spatial Fourier transform produces an electric-field image. Unlike the FFT-Correlator, EPIC can also be implemented on non-redundant arrays, including arrays where the antenna beams are non-identical. For highly redundant arrays with identical antenna beams, EPIC becomes equivalent to the FFT-correlator. Beardsley et al. (2017) propose an iterative sky-based calibration algorithm, EPICal, for such a correlator that does not require generating real-time visibility products and scales as $\mathcal{O}(N)$. However, EPICal requires prior knowledge of antenna beams that can be difficult to model or measure *in situ* at low radio frequencies. Moreover, the lack of accurate diffuse-sky models that also account for polarization at these frequencies could make it harder to decouple the sky-signal from beam models. In this paper, we choose to discuss only redundant array layouts, where the redundancy can be exploited for calibration.

An ideal calibration scheme for FFT-correlators, must be capable of minimizing the scatter in redundant visibilities because any residual scatter will become additional noise on the visibility returned by the correlator. Additionally, the calibration scheme must produce an output that can be applied to antenna voltages. Liu & Shaw (2019, section 9) summarize calibration methods that can be applied to redundant arrays. Redundant-baseline calibration (Wieringa 1992; Liu et al. 2010; Noorishad et al. 2012; Marthi & Chengalur 2014) that has been used to calibrate the Donald C. Backer Precision Array for Probing the Epoch of Reionization (PAPER; Parsons et al. 2010; Ali et al. 2015; Kolopanis et al. 2019), LOFAR (Noorishad et al. 2012), MITEoR (Zheng et al. 2014, 2017), and HERA (Dillon et al. 2020) results in complex antenna gains that can be applied to antenna voltages. The multiplicative antenna gains are computed by solving a system of equations that minimize the scatter in calibrated redundant visibilities.

A known caveat of redundant-baseline calibration is that it can only yield relative antenna gains (Liu et al. 2010; Dillon et al. 2018), i.e. the equations can constrain the ratio of antenna gains but cannot determine their actual value. The system of equations has a null space with four degenerate parameters including the absolute amplitude and the phase of antenna gains. However, this is not a problem for calibrating voltages for the purpose of FFT-correlation, since it requires only relative calibration of antennas so that visibilities of redundant baselines can be averaged coherently. Absolute calibration, to determine the degenerate parameters, can still be performed offline with the visibilities generated by the FFT-correlator.

Applications of redundant-baseline calibration, so far, had the full visibility matrix available for constructing the system of equations for calibration. However, redundant-baseline calibration does not inherently require all $N(N-1)/2$ visibilities measured at high

signal-to-noise ratio (SNR). This paper explores two redundant-baseline calibration schemes that can use $\mathcal{O}(N \log N)$ computational resources for generating visibilities for the purpose of calibration and are henceforth referred to as *reduced redundant-baseline calibration* schemes.

The FFT-correlator architecture assumed in this paper is similar to the one proposed by Zheng et al. (2014). This is shown in panel (b) of Fig. 1. The first stage, computing a spectrum of antenna voltages, is similar to the FX-correlator architecture. The yellow boxed region shows the FFT-correlation where a spatial Fourier transform on calibrated voltages results in the time-integrated unique visibilities of the array. Notice that the FFT-correlator does not produce the full visibility matrix; the redundant baselines are averaged by the spatial Fourier transform.

The blue boxed region in Fig. 1 shows the *calibrator*, which is the main focus of this paper. It performs two functions: (a) cross-correlate the baselines required for calibration in a manner similar to the second stage of an FX-correlator and (b) compute antenna gains by applying one of the two reduced redundant-calibration schemes on this set of visibilities. The computation- and resource-intensive stage of the calibrator is the first step of cross-correlating antenna pairs. The number of baselines that need to be cross-correlated in a given integration cycle determines the computational resources required by this stage of the calibrator. The reduced redundant-baseline calibration scheme employed by the calibrator dictates the set of visibilities that need to be cross-correlated and hence determines the size of the calibrator. Both the calibration schemes discussed in this paper can be adapted to a calibrator that scales as $\mathcal{O}(N \log N)$, keeping the size of the calibrator comparable to the size of the FFT-correlator.

In the first reduced redundant-baseline calibration scheme, *low-cadence calibration*, the calibrator computes the full visibility matrix by cycling through baseline pairs. As an extreme example, if the computational resources allocated to the calibrator can only cross-correlate one antenna pair at a time, the low-cadence calibrator will generate the full visibility matrix by cycling through all the antenna pairs in the array. The full visibility matrix, constructed in this fashion, is then used for redundant-baseline calibration. While all the $\sim N^2/2$ visibilities are used for redundant-baseline calibration, they are computed by using only a small fraction of the resources of an $O(N^2)$ FX-correlator. By adjusting the integration time and the number of visibilities computed within each cycle, the computational resources required by a low-cadence calibrator can be limited to an $\mathcal{O}(N \log N)$ scaling.

The second reduced redundant-baseline calibration scheme, *subset redundant calibration*,[1] is a generalization of hierarchical redundant-baseline calibration described by Zheng et al. (2014) (see the Appendix). In subset redundant calibration, the calibrator computes a partial visibility matrix by cross-correlating only a limited set of antenna pairs. Redundant-baseline calibration is applied to this partial visibility matrix to estimate the antenna gains. For example, in highly redundant arrays, since the shortest baselines involve all the antennas in the array, it is often possible to compute antenna gains by performing redundant-baseline calibration on just the shortest baselines. Since only a subset of the full visibility matrix is generated for the purpose of calibration, this technique is called subset redundant calibration. Depending on the baseline-types chosen for redundant-baseline calibration, the computational resources required by such a calibrator can also be limited to an $\mathcal{O}(N \log N)$ scaling.

[1]Shortened from subset redundant-baseline calibration.

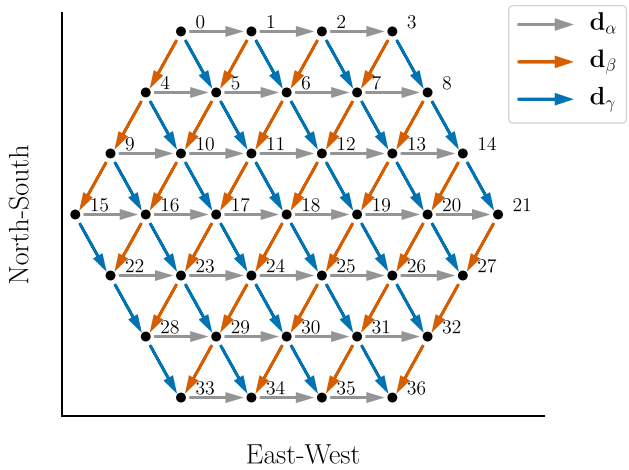

Figure 2. A hexagonal redundant array, that is loosely based on HERA, is used as a prototype for demonstrating the properties of reduced redundant-baseline calibration in this paper. Three different redundant baseline groups are marked by blue, orange, and grey arrows.

In the rest of this paper, we attempt to show that the gains estimated using either low-cadence calibration or subset redundant calibration, are capable of minimizing the scatter in the visibilities of redundant baselines. We lay out metrics for comparing the two reduced redundant-baseline calibration schemes and assessing their scalability to large-N arrays. Ultimately, we attempt to show that the calibrator design proposed here makes self-contained $\mathcal{O}(N \log N)$ correlators conceivable for future generation large-N arrays.

The layout of the rest of the paper is as follows. Section 2 quantifies the parameters that are important for understanding the performance of either reduced redundant-baseline calibration method in our simulations (Section 3). Sections 4 and 5 examine low-cadence calibration and subset redundant calibration, respectively, and discuss the limits within which they result in convergent gain solutions. Section 6 compares the performance of both methods for arrays of various sizes and discusses the limitations and advantages of employing either method for calibrating large-N arrays. Section 7 presents the conclusions of this paper.

2 METRICS TO EVALUATE REDUCED REDUNDANT-BASELINE CALIBRATION

Redundant-baseline calibration computes per-antenna complex gains by minimizing the scatter in the visibilities of redundant baselines, which makes it a suitable calibration scheme for FFT-correlators. It relies on the fact that pairs of antennas with the same beam patterns, spaced at equal distances, measure the same visibility. If V_{ij} is the visibility product of two antennas spaced a distance d apart in the east–west direction, then the visibility measured by two different antennas, V_{lm}, is the same as V_{ij} if they are also spaced a distance d apart in the same direction. In practice, this is often not true because of variations in amplifier gain, timing differences originating in the correlator, cable delays etc., that need to be calibrated. By comparing visibilities that are theoretically identical, it is possible to infer the calibration parameters for the antennas involved.

In the case of highly redundant arrays such as HERA, PAPER, CHIME, and the MWA Phase-II hexes, there are many more visibility measurements than unique baselines. This allows one to build a system of equations, which can be solved to estimate all the antenna calibration parameters. For the array layout shown in Fig. 2, the

system of equations can be constructed as[2]

$$\begin{aligned}
V_{01}^{\rm meas} &= g_0 g_1^* V_\alpha^{\rm true} + n_{01}\\
V_{12}^{\rm meas} &= g_1 g_2^* V_\alpha^{\rm true} + n_{12}\\
&\vdots \quad \text{(baselines with } \boldsymbol{d}_\alpha = \boldsymbol{r}_0 - \boldsymbol{r}_1)\\
V_{04}^{\rm meas} &= g_0 g_4^* V_\beta^{\rm true} + n_{04}\\
&\vdots \quad \text{(baselines with } \boldsymbol{d}_\beta = \boldsymbol{r}_0 - \boldsymbol{r}_4)\\
V_{05}^{\rm meas} &= g_0 g_5^* V_\gamma^{\rm true} + n_{05}\\
&\vdots \quad \text{(baselines with } \boldsymbol{d}_\gamma = \boldsymbol{r}_0 - \boldsymbol{r}_5)\\
V_{02}^{\rm meas} &= g_0 g_2^* V_\delta^{\rm true} + n_{02}\\
&\vdots\\
&\text{longer baselines,}
\end{aligned} \tag{1}$$

where $V_\alpha^{\rm true}$ is the unknown, model true visibility of all the baselines with a displacement vector $\boldsymbol{d}_\alpha$, $V_\beta^{\rm true}$ is the unknown model true visibility of baselines with displacement vector $\boldsymbol{d}_\beta$, and so on. Baselines with the same displacement vector are said to be of the same baseline-type. $V_{ij}^{\rm meas}$ is the visibility measured by the pair of antennas (i, j) in the field, and n_{ij} is the noise in that measurement. The per-antenna, complex gains denoted by g_i represent the calibration parameters of the antennas involved in measuring that visibility.

The redundant-baseline calibration process estimates the gains and models true visibilities that best describe the measured visibilities. When the full visibility matrix $\mathbf{V}^{\rm full}$ is used for the set of measured visibilities $V_{ij}^{\rm meas}$, the $V_\alpha^{\rm true}$ returned by the redundant-baseline calibration process represents the minimum-scatter average visibility for that unique baseline-type.

The system of equations in equation (1) can also be built using the visibility matrix computed by the calibrator $\mathbf{V}^{\rm reduced}$. In the case of low-cadence calibration, this set of visibilities may have a lower SNR than the full visibility matrix due to smaller integration times in the calibrator. In case of subset redundant calibration, this set of visibilities is smaller than the full visibility matrix (but sufficient to determine and over-constrain all the variables in the system of equations) due to fewer cross-correlations computed by the calibrator. The V_αs estimated by either reduced redundant-baseline calibration schemes are discarded and only the antenna gains are used to calibrate voltages for the spatial Fourier transform. The redundant-baseline averaged visibilities for all unique baseline-types are computed by the FFT-correlator.

For the purpose of this paper, it is useful to consider an intermediate hypothetical step in the spatial Fourier transform where the redundant baselines have not yet been averaged. At this stage, the visibilities in the FFT-correlator would be equivalent to the calibrated visibilities of an FX-correlator. The baseline averaging stage within the FFT-correlator, that generates $\mathbf{V}^{\rm unique}$, can be written in terms of the full visibility matrix as

$$V_\alpha^{\rm unique} = \frac{1}{N_\alpha} \sum_{(i,j)\in\alpha} \frac{V_{ij}^{\rm full}}{,} g_i g_j^* \tag{2}$$

[2]In general, all the measurements and variables in this system of equations have a time and frequency dependence. We have omitted writing this explicitly for notational convenience.

where N_α is the number of redundant baselines that contribute to that baseline-type. Since it is easier to quantify the effect of calibration on visibilities than on voltages, we use this equation to represent the process of calibration in the FFT-correlator.

Another difference between traditional redundant-baseline calibration and reduced redundant-baseline calibration, that is evident from equation (2), is that the latter estimates antenna gains from a different set of visibilities ($\mathbf{V}^{\rm reduced}$) than what they are finally applied to ($\mathbf{V}^{\rm full}$). In this section, we discuss two metrics that will help in evaluating the effect of this: (a) the uncertainty in estimated antenna gains and (b) scatter in visibilities calibrated by a reduced redundant-baseline calibration process.

The uncertainty in antenna gains has to decrease or remain constant with increase in array size, for the reduced redundant-baseline calibration scheme, and consequently the calibrator design, to be scalable to large arrays. Moreover, the uncertainty in the estimated gains can be expressed in terms of the SNR of the measured visibilities and the number of baselines used in the calibration process, providing us with a convenient metric to directly compare low-cadence calibration and subset redundant calibration. The overall uncertainty in the redundant-baseline averaged calibrated visibilities comes from both the noise in the measured visibilities and the uncertainty in the estimated gains. A quantitative measure of the antenna gain uncertainty will help us estimate the contribution of gain errors to the overall error in the calibrated visibilities.

The scatter in visibilities post-calibration is a direct probe of the effectiveness of the reduced redundant-baseline calibration process in estimating gains that can calibrate antenna voltages for the FFT-correlation. The spatial Fourier transform averages the visibilities of redundant baselines (equation 2) converting any residual scatter into noise in the estimated visibilities. Hence, quantifying the post-calibration scatter will help us evaluate the performance of either reduced redundant-baseline calibration scheme with respect to traditional redundant-baseline calibration.

An important mathematical detail, before delving into the metrics that assess reduced redundant-baseline calibration schemes, is that the system of equations represented by equation (1) is not linear. Wieringa (1992) suggests a logarithmic approach to linearizing which can be written as

$$\ln V_{ij}^{\rm meas} = \ln g_i + \ln g_j^* + \ln V_\alpha^{\rm true} + n'_{ij}. \tag{3}$$

The noise parameter n'_{ij} evaluates to a non-Gaussian error that depends on the SNR of the measured visibilities. Liu et al. (2010) discuss the noise-bias in antenna gains due to this non-Gaussianity, and propose another approach based on Taylor expanding the variables around a starting point. This paper employs a widely used third approach, called `omnical` (Zheng et al. 2014; Ali et al. 2015; Li et al. 2018; Dillon et al. 2020), that was originally developed for the MITEoR experiment. We use the logarithmic approach to make theoretical arguments about the nature of gain solutions, since constant coefficients make the system of equations easier to analyse. However, simulations run to test these arguments and the plots shown in this paper have been generated by employing the `omnical` algorithm that is not noise-biased. In general, most of the results presented in this paper are not dependent on the solving technique used.

2.1 Uncertainty in antenna gains

The uncertainty in the antenna gains, estimated by solving a system of equations, is given by the variance–covariance matrix (covariance matrix henceforth). The linearized system of equations shown in

equation (3) can be written in matrix notation as $\mathbf{Ax} + \mathbf{n} = \mathbf{b}$, where $\mathbf{A}$ is a constant complex-valued matrix of dimensions $N_m \times N_v$ (N_m is the number of measured visibilities and N_v is the number of variables). $\mathbf{x}$, $\mathbf{b}$, and $\mathbf{n}'$ are one-dimensional matrices of the variables (log-gains and log-unique-visibilities), measured quantities (log-visibilities), and noise in each measurement, respectively. The covariance matrix $\mathbf{C}$ (of dimensions $N_v \times N_v$) for the estimated variables is given by

$$\mathbf{C} = \langle \mathbf{x}\mathbf{x}^\dagger \rangle = (\mathbf{A}^\dagger \mathbf{N}^{-1} \mathbf{A})^{-1}. \tag{4}$$

The diagonal of the covariance matrix gives the variance in estimated variables, including antenna gains. The noise covariance matrix ($\mathbf{N} = \langle \mathbf{n}\mathbf{n}^\dagger \rangle$) is a statistical estimate of the noise in the measurement matrix $\mathbf{b}$. The matrix $\mathbf{n}$ cannot be measured and can only be estimated from the thermal noise expected in the measurements.

The covariance matrix $\mathbf{C}$ reflects the covariance between all variables returned by the reduced redundant-baseline calibration process. Since the calibrator uses only the antenna gains and discards the model visibilities estimated by redundant-baseline calibration, the covariance matrix of interest is the *marginalized* covariance ($\mathbf{C}'$) of just the antenna gain solutions. Assuming all variables are normally distributed, the marginalized covariance matrix for a subset of variables is given by the rows and columns of the variables of interest. Hence, the marginalized covariance matrix of the gain solutions is given by the first N rows and columns of the covariance matrix in equation (4), where N is the number of antennas in the array.

Liu et al. (2010, section 2.4) derive the noise covariance matrix for the logarithmic approach to linearizing (equation 3) under the assumptions that the measured visibilities have a high SNR, and that the noise in the measured visibilities is uncorrelated between baselines, Gaussian in nature and similar across all baseline-types. This noise covariance matrix evaluates to

$$\mathbf{N} \approx (\mathrm{SNR})^{-2}\, \mathbf{I} \tag{5}$$

under the additional assumption that all the baselines in the array have the same average SNR. In general, this assumption does not hold when observing a real sky. However, in this paper, we only use SNR in the context of other array-averaged parameters, for which this assumption is justified. Since we have assumed that the noise is uncorrelated between baselines, the noise covariance matrix is just proportional to the identity matrix $\mathbf{I}$.

Substituting the noise covariance matrix into equation (4), and taking the first N rows and columns of the covariance matrix, which we denote by $(\mathbf{A}^\dagger \mathbf{A})^{-1}_{(N\times N)}$, we get the covariance of antenna gains. The diagonal of this matrix represents the variance or uncertainty in the estimated antenna gains:

$$\sigma_g^2 \approx (\mathrm{SNR})^{-2}\, \mathrm{diag}\left[(\mathbf{A}^\dagger \mathbf{A})^{-1}_{(\mathrm{N}\times\mathrm{N})}\right]. \tag{6}$$

The two reduced redundant-baseline calibration schemes effect the uncertainty in gains according to the above equation. In low-cadence calibration, the lower SNR of visibilities in the calibrator result in higher gain variance as compared to traditional redundant-baseline calibration. In subset redundant calibration, only a sub-matrix of $\mathbf{A}$ is used for calibration, again resulting in antenna gains with a higher variance.

The matrix $(\mathbf{A}^\dagger \mathbf{A})$ is nearly diagonal (small off-diagonal terms), with each entry equal to the number of equations in which the corresponding variable is involved. When redundant-baseline calibration is performed using the full visibility matrix, the first N diagonal entries of this matrix are each equal to N since every antenna is involved in N equations. At constant SNR, this results in the following scaling for gain variance:

$$\sigma_g^2 \propto \frac{1}{N}. \tag{7}$$

In Sections 4 and 5, we derive a scaling relation for the variance in gains estimated using a $\mathcal{O}(N \log N)$ calibrator that implements low-cadence calibration and subset redundant calibration respectively, and compare it with the above relation.

2.2 Scatter in visibilities of redundant baselines

The spatial Fourier transform in an FFT-correlator, which decreases the computational scaling from $\mathcal{O}(N^2)$ to $\mathcal{O}(N \log N)$, averages the visibilities of redundant baselines. Hence, the post-calibration residual scatter in redundant visibilities is an important metric for assessing the gains estimated by a reduced redundant-baseline calibration process.

The scatter in redundant visibilities is quantified by the reduced χ^2 of antenna gains and model visibilities (χ_r^2) which is given by

$$\chi_r^2 = \frac{1}{\mathrm{DoF}} \sum_{(i,j)\in\alpha,\forall\alpha} \frac{\left|V_{ij}^{\mathrm{meas}} - g_i g_j^* V_\alpha^{\mathrm{true}}\right|^2}{\sigma_{ij}^2}, \tag{8}$$

where σ_{ij}^2 is the variance of the noise in measured visibilities, n_{ij}, in equation (1). DoF is the degrees of freedom in this system of equations given by

$$\mathrm{DoF} = \mathrm{N_{obs}} - \mathrm{N} - \mathrm{N_{ubl}} + 2. \tag{9}$$

N_{obs} is the total number of cross-correlations computed (or number of visibility measurements), N is the number of antennas in the array, and N_{ubl} is the number of unique baseline-types in the system of equations. The additional offset by two accounts for the number of degenerate parameters in the system of equations representing a single-polarization (Dillon et al. 2018, 2020).

The reduced redundant-baseline calibration process within the calibrator is setup to minimize the χ_r^2 between the visibility matrix computed by the calibrator $\mathbf{V}^{\mathrm{reduced}}$ and the estimated gains and model true visibilities. In an array with identical antennas and perfect redundancy, we expect this χ_r^2 to evaluate to one.

The estimated gains, however, apply to antenna voltages prior to the spatial Fourier transform that computes a different visibility matrix $\mathbf{V}^{\mathrm{unique}}$ than that used in the reduced redundant-baseline calibration process. Hence, the χ_r^2 evaluated using the gains estimated by the reduced redundant-baseline calibration process, model visibilities computed by the FFT-correlator (estimated using equation 2) against the V_{ij}^{meas} drawn from the full visibility matrix computed by an FX-correlator, is a useful a metric to assess the effectiveness of the estimated gains in calibrating the whole array.

In Section 6, we use the χ_r^2 computed in this way to compare the performance of reduced redundant-baseline calibration to traditional redundant-baseline calibration with the full visibility matrix.

3 SIMULATION

In the following sections, we discuss the performance of low-cadence calibration and subset redundant calibration using simulated visibilities and antenna gains. We use hexagonal array layouts, like the one shown in Fig. 2, with varying number of antennas for the simulations. Though this layout is loosely based on HERA (see Dillon & Parsons 2016), we expect the derived trends to hold for any two-dimensional redundant array layout. We assume perfect redundancy in the array and identical antenna beams.

The simulations start by generating a set of true visibilities for all the unique baseline-types and a set of gains for all the antennas in the array. These simulated true visibilities have a random constant average amplitude across all baseline-types are constant in time and uncorrelated between baseline-types. While this does not reflect a real sky signal, it is sufficient for the purpose of this paper because redundant-baseline calibration only has a weak dependence on the actual signal observed. The simulated antenna gains are Gaussian-distributed, with an average amplitude of 1. The antenna-to-antenna variation, or gain scatter, is assumed to be at the ∼10 per cent level which was found to be typical for HERA antennas (Kern et al. 2020).

The process of calibration in the FFT-correlator (the first step in the yellow boxed region of Fig. 1) is simulated by applying the estimated gains to the full visibility matrix. This is equivalent to applying the gains to individual antenna voltages and cross-correlating them. The full visibility matrix computed by an FX-correlator, $\boldsymbol{V}^{\rm full}$, is generated by multiplying the simulated gains and true visibilities and adding Gaussian random noise. In the simulations where an explicit SNR for the visibilities is not mentioned, we have assumed an SNR of 10 even though a more favourable SNR is expected in practice. The visibilities produced by the FFT-correlator $\mathbf{V}^{\rm unique}$ are generated by applying the gains estimated by one of the reduced redundant-baseline calibration processes to this full visibility matrix (equation 2).

For low-cadence calibration, the visibility matrix computed within the calibrator, $\mathbf{V}^{\rm reduced}$, is generated by adding higher amplitude Gaussian noise to the multiplied gains and true visibilities. For subset redundant calibration, this visibility matrix is generated by choosing the visibilities of the selected antenna pairs from the full visibility matrix. Reduced redundant-baseline calibration is performed on this visibility matrix using the `omnical` algorithm, with a damping factor of 0.3 and convergence criteria of 10^{-10}. The amplitude and phase degeneracies of the resulting gains are fixed by comparing with the amplitude and phase of the simulated input gains. Unless otherwise specified, the variance of antenna gains is computed by running 256 simulations with the same underlying gains and different realizations of true visibilities.

4 LOW-CADENCE CALIBRATION

Low-cadence calibration is a reduced redundant-baseline calibration scheme that estimates antenna gains from visibilities that have been computed in a round-robin fashion. The calibrator in Fig. 1 cross-correlates the baselines required for calibration, but the computational resources allocated to it cannot scale faster than $\mathcal{O}(N \log N)$. The computational resources required to compute visibilities is determined by the number of baselines that need to be cross-correlated *simultaneously*. By decreasing the number of antenna pairs that need to be correlated at a time, the computational resources required by the calibrator can be reduced. The full visibility matrix is populated after a few cycles and this is used to redundantly calibrate the array. Since redundant-baseline calibration can only be performed once in a given number of cross-correlation cycles, this calibration scheme is called low-cadence calibration.

The size of a low-cadence calibrator is determined by two parameters – the time period available for generating the full visibility matrix within the calibrator ($t_{\rm cal}$) and the integration time allotted to each cycle ($t_{\rm int}$). These are related by the equation:

$$t_{\rm cal} = N_{\rm cycle} \times t_{\rm int} \tag{10}$$

where $N_{\rm cycle}$ is the number of integration cycles taken by the calibrator to populate the full visibility matrix. The size of the low-cadence calibrator is inversely proportional to $N_{\rm cycle}$ i.e. for a small calibrator size we require a large $t_{\rm cal}$ and a small $t_{\rm int}$.

Redundant-baseline calibration operates under the assumption that antenna gains and true visibilities are constant during the time period required to compute all the visibilities involved in the system of equations. For an FX-correlator, this is equal to the integration time of the full visibility matrix which is usually smaller than the inherent gain variability, and necessarily smaller than the time period over which the visibilities evolve due to a constantly rotating sky. For a low-cadence calibrator, the constancy of antenna gains and true visibilities within a calibration cycle has to be manually enforced.

The upper limit of $t_{\rm cal}$ is set by the inherent gain variability of the array, which could depend on numerous factors including the analogue signal chain, the radio frequency interference environment and the precision of antenna gains required for the science application. If the time taken to generate the full visibility matrix is larger than the interval within which gains can be assumed to be constant, redundant-baseline calibration can result in erroneous gain solutions. If the time period of gain variability is large, it is possible that the true visibilities change within this period. However, to preserve redundancy we only require that all pairs of antennas with the same baseline be correlated simultaneously. Since this is necessarily always less than N visibilities, a calibrator which can cross-correlate at least N baselines can accommodate the largest redundant-baseline group in the array and satisfy this condition.

A realistic lower limit for $t_{\rm cal}$ is the integration time of visibilities in the FFT-correlator. Within this period, the assumption of constant antenna gains and true visibilities holds and redundant-baseline calibration can be solved using the algorithms currently available. While $t_{\rm cal}$ can theoretically be set to a smaller value, it could unnecessarily increase the size of the calibrator by decreasing $N_{\rm cycle}$ for a given $t_{\rm int}$.

4.1 Scaling in gain variance with integration time

The relationship between integration time and SNR of a measured visibility is given by the radiometer equation (see Thompson, Moran & Swenson 2017, Appendix 1.1). Substituting the radiometer equation into the variance of antenna gains in equation (6), we get

$$\begin{aligned}\sigma_g^2 &\approx ({\rm SNR})^{-2}\ {\rm diag}\left[(\mathbf{A}^\dagger\mathbf{A})^{-1}_{(N\times N)}\right]\\ &\propto (\sqrt{t_{\rm int}})^{-2}\ {\rm diag}\left[(\mathbf{A}^\dagger\mathbf{A})^{-1}_{(N\times N)}\right]\ \propto \frac{1}{t_{\rm int}},\end{aligned} \tag{11}$$

which quantifies the variance of the antenna gains estimated by performing redundant-baseline calibration on visibilities integrated for a given duration. A shorter integration time leads to lower SNR in the measured visibilities and consequently, a higher variance in the antenna gains.

Fig. 3 shows the trend in estimated gain variance with the average SNR of all the visibilities in an array. Each box in the figure represents the distribution of gain variance over all antennas in the array and has been generated assuming that all the baseline-types in the array have the same specified SNR. When the average SNR of visibilities is high, the estimated antenna gain variance follows the expected inverse square relationship. At low SNR, when the theoretical gain variance estimated using equation (6) becomes comparable to the antenna-to-antenna scatter in gains squared (∼0.01 for this simulation), redundant-baseline calibration fails at estimating antenna gains. Below a threshold SNR, that is set by the gain scatter, the inverse square relationship breaks down and the variance of estimated antenna gains becomes dependent on the solver. That is, the logarithmic approach, Taylor expansion approach, and `omnical`

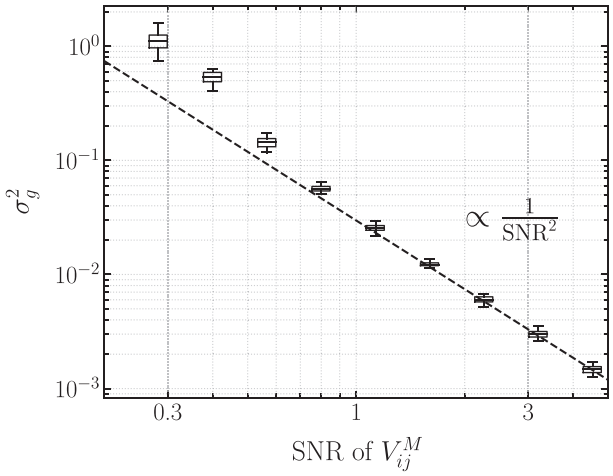

Figure 3. Relationship between the variance of antenna gains and the SNR of visibilities used for low-cadence calibration. The dashed line represents the inverse square relationship predicted by equation (6) and the boxes show the results of simulation. The estimated variance of antenna gains in each simulation is a distribution of N points, where N is the number of antennas in the array. The upper and lower limits of the 'box' represent the upper and lower quartiles of the distribution. The horizontal bar within the box represents the mean of the distribution, and the whiskers show the total range of the estimated gain variances. At high SNR, the gain variance follows the theoretically expected trend. At low SNR, the variance is higher than the predicted value and becomes solver dependent because the χ^2 is not effectively minimized by the solver.

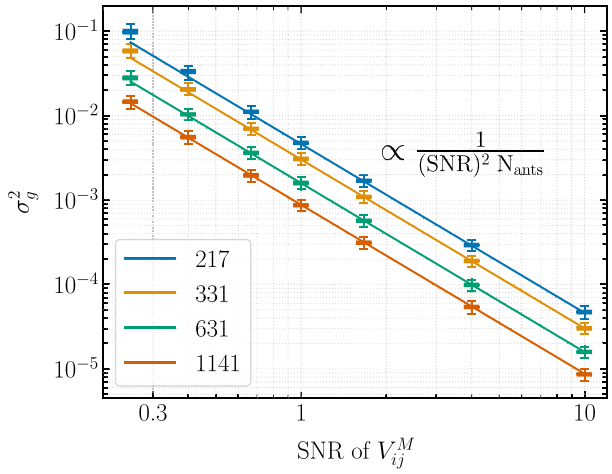

Figure 4. Relationship between distribution of gain variance and SNR of visibilities used in low-cadence calibration. The four different colours represent hexagonal arrays of various sizes; legend shows the number of antennas in each layout. The solid lines represent the theoretical trend predicted by equation (6) and the boxes show the distribution of gain variance obtained from simulation (see Fig. 3 for explanation). The threshold SNR, below which the variance deviates from the theoretical trend, is dependent on the number of antennas in the array.

algorithms of linearizing the system of equations result in different deviations from the given trend. This is because each algorithm minimizes χ^2 in a different way and none of them are effective at converging to the solution.

4.2 Scaling in gain variance with an $\mathcal{O}(N \log N)$ calibrator

Fig. 4 shows the trend in gain variance with SNR of measured visibilities for hexagonal layouts with different number of antennas. The variance is suppressed by a factor of N as the number of antennas in the array increase, because of the $(\mathbf{A}^\dagger\mathbf{A})^{-1}_{(\mathrm{N}\times\mathrm{N})}$ term in equation (6). Notice that in larger arrays, the gain variance follows the theoretical trend even when the SNR of measured visibilities is less than one.

The threshold SNR below which the gain variance diverges from the theoretical prediction, changes with the number of antennas in the array. When the theoretical gain variance is less than the square of the expected gain scatter, the gain variance follows the inverse square relationship *even when the SNR of measured visibilities is less than one*. This result is important because it allows low-cadence calibration to be scaled to extremely large arrays.

Say, the computational resources allocated to the calibrator are restricted to scale similarly as the FFT-correlator. The calibrator cross-correlates $pN\log N$ baselines in each integration cycle, where p is a pre-factor (like a proportionality constant) to convert the $\mathcal{O}(N \log N)$ scaling into number of baselines. A larger pre-factor would result in a larger calibrator size. For a fixed interval of calibration t_{cal}, the integration time is smaller for larger arrays according to the scaling:

$$t_{\mathrm{int}} = \left(\frac{pN\log N}{N^2/2}\right) t_{\mathrm{cal}} = \left(\frac{2p\log N}{N}\right) t_{\mathrm{cal}}. \tag{12}$$

Hence, even if the interval of calibration is large, the multiplying factor might become small enough to push the SNR of measured visibilities to less than one. Substituting the result of equation (12) into equation (11), we get

$$\sigma_g^2 \propto \left(\frac{N}{2p\log N}\cdot\frac{1}{t_{\mathrm{cal}}}\right)\frac{1}{N} \propto \left(\frac{1}{p\log N}\right)\cdot\frac{1}{t_{\mathrm{cal}}}, \tag{13}$$

which shows that the antenna gain variance improves with array size even at constant t_{cal}. That is, even though the SNR of measured visibilities might decrease to a value less than one, the antenna gain variance decreases. The price that one pays for not using $\mathcal{O}(N^2)$ resources for calibration is that the precision in antenna gains scales more slowly as compared to that of redundant-baseline calibration with the full visibility matrix which is given by equation (7).

Low-cadence calibration is a calculated way of trading computational resources for precision in the antenna gain solutions. As long as the size of the calibrator scales faster than $\mathcal{O}(N)$ with the size of the array, the variance in antenna gains decreases with increase in the number of antennas. A potential drawback of low-cadence calibration, especially when applied to arrays with over 10 000 antennas, is the time taken by a linearized solver to result in convergent gains. Dillon et al. (2020) show that the time taken by the `omnical` algorithm scales as $\mathcal{O}(N^2)$ when the solver has to optimize N^2 baselines. If the time interval between calibration cycles t_{cal} can be proportionally decreased, it might still be possible to obtain real-time solutions. However, t_{cal} is usually set by the inherent gain variability in the array which might not be scalable with array size. One way of addressing this issue is to look at redundant-baseline calibration with a limited set of baselines.

5 SUBSET REDUNDANT CALIBRATION

The spatial Fourier transform in an FFT-correlator averages visibilities of redundant baselines. Traditional redundant-baseline calibration assumes that all the $\sim N^2/2$ cross-correlation products are available for calibration, which could be non-viable to compute for

large-N arrays. Subset redundant calibration is a reduced redundant-baseline calibration scheme that attempts to estimate antenna gains from visibilities of only a limited set of baseline-types. This section examines the effect of not using all baselines for redundant-baseline calibration on the variance of estimated antenna gains.

While considering baselines for subset redundant calibration, it is useful to distinguish between baseline-types (or unique baselines) and redundant baseline groups. A baseline-type that a particular antenna pair belongs to is specified by the displacement vector between the two antennas. A redundant baseline group consists of all the antenna-pairs that have the same baseline vector. For instance, Fig. 2 shows three different baseline-types with displacement vectors pointing east (grey), south-east (blue) and Southwest (orange). Each baseline-type has 30 different antenna pairs in its redundant baseline group, marked in arrows of the same colour.

5.1 Brief discussion on using short baselines

The short baseline-types are important in subset redundant calibration for two main reasons: (a) they involve all the antennas in the array, allowing redundant-baseline calibration on just these visibilities to estimate all the antenna gains and (b) the redundant baseline groups of these baseline-types are larger than the longer baseline-types. For example, in the layout shown in Fig. 2, there are 30 baselines in each of the redundant baseline groups that belong to the shortest three baseline-types, while there are only four baselines that belong to the group formed by the baseline-type like (0,31). This is important for subset redundant calibration because every new baseline-type added to the system of equations requires a new variable in the form of the unique visibility for that baseline-type. Since the short baselines have a higher ratio of redundant baselines (measurements) to unique visibilities (variables), they contribute more to constraining the gain solutions.

In addition to this, at low radio frequencies, the short baselines pick up the bright diffuse emission from our galaxy and have high SNR visibilities. As shown in equation (6), this suppresses gain variance and results in higher precision gain solutions. Orosz et al. (2018) discuss other advantages of using only short baselines from the perspective of non-redundancies in a realistic array layout. They argue that calibration errors affect the inferred power spectrum, and contamination worsens when longer baselines are included in the redundant-baseline calibration process. Li et al. (2018) point out that redundant-baseline calibration performs better than sky-based calibration at low radio frequencies, partly because short baselines have to be ignored for sky-based calibration due to poor diffuse sky models. On the other hand, the shorter baselines are more susceptible to systematic errors like antenna cross-coupling (Kern et al. 2019) and may be more non-redundant than the longer baselines (Dillon et al. 2020) in a realistic array layout.

A practical subset redundant calibrator would cross-correlate a combination of short and long baselines that produces the best estimate of antenna gains for the array. Since the voltages from all the antennas are available to the subset redundant calibrator, the combination of baselines that it needs to compute can also dynamically change with time/day of observation. In this paper, the baselines used to perform subset redundant calibration are considered in the order of baseline length from shortest to longest. That is, a smaller calibrator preferentially cross-correlates only the shorter baselines. However, the results presented in this section apply to combinations of short and long baselines as well.

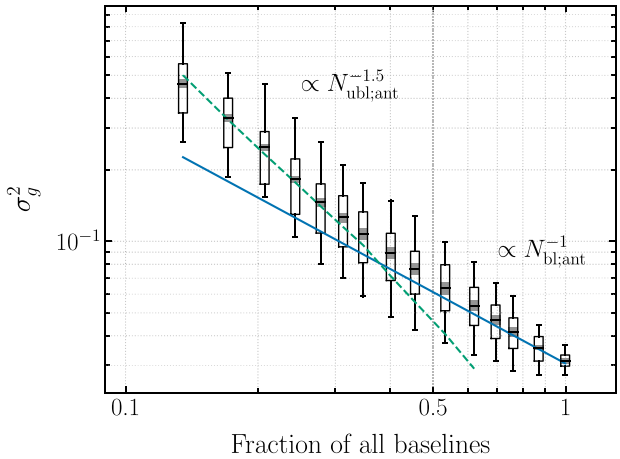

Figure 5. Relationship between fraction of baselines used in subset redundant calibration and the resulting variance in estimated gains. The boxes show the distribution of gain variances in simulation (see Fig. 3 for explanation), where increasingly longer baselines are included in the subset redundant calibration system of the simulation. When the fraction of baselines used in estimating antenna gains is small, the gain variance depends on the number of unique visibilities per antenna used in calibration ($N_{\mathrm{ubl;\,ant}}$); this trend is shown by the dashed green line. When a large fraction of baselines are used for calibration, the gain variance depends on the total number of measurements per antenna ($N_{\mathrm{bl;\,ant}}$); shown by the solid blue line. A second set of boxes in grey show the antenna-to-antenna variation with all the edge antennas excluded.

5.2 Degeneracy criterion

When using a limited set of baselines for redundant-baseline calibration, it is important to ensure that there are sufficient number of measurements to determine all the variables. The solution space of equation (1) has a null space with following four degenerate parameters (Wieringa 1992; Liu et al. 2010; Dillon et al. 2018):

(a) The absolute amplitude of the gains (or the sum of all gains)
(b) The absolute phase of the gains (or the sum of all gain phases)
(c) The phase slope of the gains in the x-direction
(d) The phase slope of the gains in the y-direction

When selecting baselines for subset redundant calibration, it is important to verify that the null space of the solution set is restricted to these four degenerate parameters. Introducing more degeneracies allows the gain solutions to vary in that dimension and could require future corrections. If the additional degeneracy does not have a physical interpretation, this may not even be possible. For a hexagonal layout like the one shown in Fig. 2, a minimum of three unique baseline-types, with displacement vectors in the directions marked in the figure, are required to satisfy the degeneracy requirement.

5.3 Scaling in gain variance with number of baselines

Fig. 5 shows the relationship between number of baselines used in subset redundant calibration and the corresponding variance in the estimated gains. The number of baselines are shown in terms of a fraction of the total baselines in the array. The black boxes show the results of simulation (see Section 3). The x-axis from left to right represents baselines added to the subset redundant calibration system in ascending order of baseline length (starting with the minimum required to satisfy the degeneracy criterion). For each baseline-type added to the system, it is assumed that all the redundant baseline pairs that contribute to that baseline-type are used for calibration.

The SNR of visibilities is assumed to be similar for all baseline-types (unlike for a real sky) and constant through the simulation. The vertical range of the boxes, which represents distribution of gain variance over antennas in the array, is larger than the case of low-cadence calibration. However, most of this variation comes from the edge antennas in the array. A second set of solid boxes in grey shows the distribution of gain variances with the edge antennas excluded. The reason for this decrease in variation is discussed in more detail in Section 5.4.

The solid line, in blue, represents the inverse of the total number of baselines per antenna ($N_{\mathrm{bl;\,ant}}$) that are used to perform subset redundant calibration:

$$N_{\mathrm{bl;ant}} = \frac{N_{\mathrm{obs}}}{N} \approx \frac{f\ N^2/2}{N} = f\ \frac{N}{2}. \tag{14}$$

Here, f represents the fraction of all baselines used in subset redundant calibration. When the fraction of baselines used in subset redundant calibration is high, the gain variance asymptotes to the inverse of the total number of measurements-per-antenna in the system of equations. This trend is expected from the Gaussian noise in visibility measurements; each new measurement added to the system contributes to decreasing the noise in the estimated gains.

When the fraction of baselines used in estimating antenna gains is small, gain variance depends on two factors: (a) the number of baseline-types included in subset redundant calibration and (b) the number of antennas that are involved in forming redundant baselines for these baseline-types. A combination of these two factors is captured by the variable $N_{\mathrm{ubl;\,ant}}$ that represents the average number of unique visibilities per antenna:

$$N_{\mathrm{ubl;ant}} = \frac{1}{N}\sum_{\alpha} N_{\mathrm{ant}}[\in V_\alpha]. \tag{15}$$

In the above equation, the variable being summed is the number of antennas that are involved in forming redundant baselines of the baseline-type V_α. The summation runs over all the baseline-types that are used in subset redundant calibration and N is the total number of antennas in the array. Consider the case where all the baseline-types used in subset redundant calibration have redundant baselines involving all the antennas in the array (for instance, when only the shortest three to six baseline-types are used for calibration). In such a system, $N_{\mathrm{ubl;\,ant}}$ simply evaluates to the total number of baseline-types (or unique visibilities) used in the calibration process. If some of the baseline-types used in subset redundant calibration involve only a couple of antennas, $N_{\mathrm{ubl;\,ant}}$ is smaller than the total number of unique visibilities in the system of equations.

Empirically, we find that the relationship between gain variance and the average number of unique visibilities per antenna is a power law with a slope around -1.5 for hexagonal and square layouts. This power-law trend is shown by the dashed green line in Fig. 5. The large antenna-to-antenna variation in this regime makes it difficult to determine the exact slope or understand the origin of this power law. We suspect that it originates in the way gain error propagates from antenna to antenna.

In summary, for a subset redundant calibration system that uses only a small fraction of the total baselines in the array, gain variance improves when baseline-types with larger redundant groups are used for calibration.

5.4 Covariance in estimated gains

The improvement in gain variance obtained when using a higher number of unique baselines per antenna can also be explained through the gain covariance. When the fraction of baselines used in calibration is small, in addition to high variance, the gains also have a relatively high covariance. The covariance in gains is given by the off-diagonal terms of the marginalized covariance matrix:

$$\mathbf{C}' \approx (\mathrm{SNR})^{-2}(\mathbf{A}^\dagger\mathbf{A})^{-1}_{(N\times N)}. \tag{16}$$

Both the diagonal and off-diagonal terms of the matrix $(\mathbf{A}^\dagger\mathbf{A})$ change when a subset of baseline-types are used for redundant calibration. Inverting this matrix changes the covariance in the resulting gains.

Fig. 6 shows the marginalized covariance matrices for three different subset redundant calibration systems (for the hexagonal array layout in Fig. 2). The logarithmic approach to linearizing, shown in equation (3), naturally results in separating the amplitude and phase of gains into the real and imaginary parts of the logarithm, respectively. Hence, using the real (imaginary) part of the matrix $\mathbf{A}$ gives the gain covariance in the amplitude (phase) of gains. All the covariance matrices are normalized by the thermal noise in the visibilities used for redundant-baseline calibration.

Panel (a) of Fig. 6 shows the covariance matrices for the case where redundant-baseline calibration is performed using the full visibility matrix. The variance in antenna gains, given by the diagonal of the matrix, has an average value of $1/N$ as predicted by equation (7). Though not evident in the figure, this variance is weakly dependent

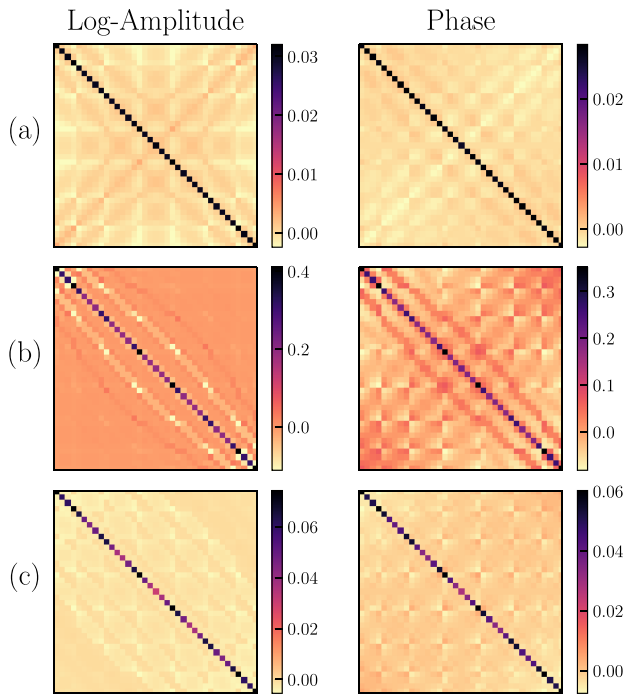

Figure 6. Covariance in the amplitude and phase of antenna gains for subset redundant calibration performed on the hexagonal layout of Fig. 2. Panel (a) shows the marginalized covariance for the case where redundant-baseline calibration is performed using the full visibility matrix. Panel (b) shows the covariance matrix when subset redundant calibration is performed with just the shortest three baselines required to satisfy the degeneracy criterion. In addition to higher gain variance, the antennas have also have a high covariance. Panel (c) is the covariance matrix for subset redundant calibration performed using more than half of the total baselines in the array. Both the variance and the covariance are comparable to that of redundant-baseline calibration performed with the full visibility matrix even though the covariance has a different structure compared to panel (a), and the variance is still dependent on antenna location.

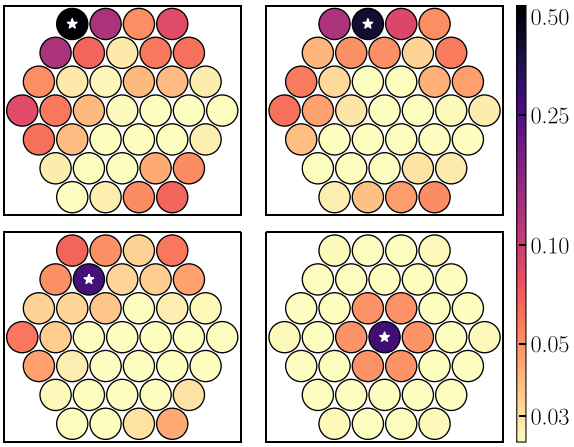

Figure 7. Four rows of the covariance matrix in Panel (b) of Fig. 6. The colourbar has been log-normalized to make the covariance more evident. The covariance shown in each panel is the sum of covariance in amplitude and phase for the antenna that is marked with a star. For all antennas, the covariance is highest with adjacent antennas because only those visibility measurements are used for constraining gain solutions. The variance and covariance of corner and edge antennas (upper two panels) are higher than that of antennas placed centrally (lower two panels) because the edge antennas participate in fewer visibility measurements than central antennas.

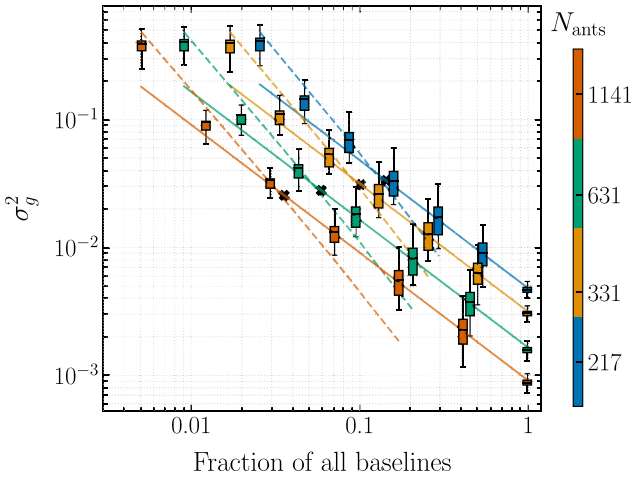

Figure 8. Relationship between variance of antenna gains and the fraction of baselines used in subset redundant calibration, for hexagonal array layouts of various sizes (shown in different colours). The boxes represent the distribution of gain variance obtained from simulation (see Fig. 3 for explanation). The solid lines represent an inverse trend of baselines per antenna ($N_{\rm bl;ant}^{-1}$). The dashed lines show a power-law dependence on the number of different baseline-types per antenna ($N_{\rm ubl;ant}^{-1.5}$) used for calibration. The small black crosses represent the fraction of baselines that can be cross-correlated by a calibrator that can process $N\log N$ baselines.

on antenna location as shown by Dillon & Parsons (2016). There a small but non-zero covariance in the antenna gains.

Panel (b) of Fig. 6 shows the covariance matrices for the extreme case where only the shortest three baselines (the minimum baseline-types required to satisfy the degeneracy criterion) are used in redundant-baseline calibration. Notice that the variance is nearly an order of magnitude higher than the first case and clearly dependent on antenna location. The covariance between antennas is non-negligible and higher between antennas that have a high variance.

Panel (c) of Fig. 6 shows the covariance matrices for the case where more than half the baselines are used in subset redundant calibration. This is the threshold at which gain variance starts following the inverse measurements per antenna trend in Fig. 5. Even though the variance is still antenna location dependent and the covariance has a different structure, they are comparable to the case of redundant-baseline calibration with the full visibility matrix.

When the subset redundant calibration system involves a larger number of unique baselines, there are more independent constraints on the gain of each antenna. This decreases the average gain variance, covariance between antennas and also the antenna-location dependence of the variance. In Fig. 5, the errorbars associated with the gain variance estimated in simulation represent the antenna-to-antenna variation within the simulation. Hence, the errorbars are larger for the gains estimated using a smaller fraction of baselines.

This dependence of antenna gain variance and covariance on the number of independent constraints per antenna is more evident in Fig. 7, which shows four different rows of the covariance matrix in Panel (b) of Fig. 6 in an exaggerated manner. In this case, subset redundant calibration is performed with just the shortest three baseline-types required to satisfy the degeneracy criterion. The covariance and variance of edge antennas is higher than that of centrally placed antennas because the corner antennas participate in fewer cross-correlations (three for the antenna in the top left-hand panel, four for the antenna in the top right) than centrally placed antennas (six each for the antennas in the bottom two panels). Hence, there are fewer independent constraints for the edge antennas, leading to a higher variance in their estimated gains. Ignoring the edge antennas after subset redundant calibration leads to an 80–30 per cent decrease in the antenna-to-antenna variation depending on the fraction of baselines used in the calibration process. The distribution of antenna gain variances with the edge antennas excluded is shown using solid grey boxes in Fig. 5.

5.5 Scaling in gain variance with an $\mathcal{O}(N\log N)$ calibrator

Fig. 8 extends the relationship between gain variance and fraction of baselines used in subset redundant calibration to hexagonal layouts with larger number of antennas. The four different colours represent four different array sizes. Within each colour, the boxes represent the distribution of gain variance over antennas in the array (see Section 3). The solid lines represent an the inverse measurements per antenna trend ($N_{\rm bl;ant}^{-1}$) and the dashed lines show a power-law relationship between gain variance and number of unique visibilities per antenna ($N_{\rm ubl;ant}^{-1.5}$). It is evident that the two asymptotes to the gain variance, shown in Fig. 5, hold with changing array size.

When a large fraction of baselines are used for subset redundant calibration, the gain variance depends on the total number of baselines used for calibration. At a constant fraction of baselines, the number of baselines formed by an antenna is proportional to the number of antennas in the array as shown in equation (14). The four solid lines show a scaling in gain variance by this factor. When the fraction of baselines used for calibration is small, the gain variance does not necessarily decrease with increase in array size. The leftmost point within each array size represents the variance in the estimated gains when subset redundant calibration is performed, using just the shortest three baseline-types (required for the degeneracy criterion). As evident from the figure, the gain variance is nearly constant despite the larger number of redundant baselines in an array with more antennas. This is because, within increasing array size, there are as many new variables (in the form of antenna gains) as measurements. For improvement in variance, the size of the subset redundant calibrator should increase with the array size.

Say the computational resources allocated to a subset redundant calibrator are restricted to scale similarly to the FFT-correlator. The calibrator can process $pN\log N$ baselines for a given array size. The fraction of baselines that can be processed by such a calibrator is given by

$$f = \frac{pN\log N}{N^2/2} = \frac{2p\log N}{N}. \tag{17}$$

Fig. 8 shows this fraction of baselines, for a pre-factor of one, in small black crosses. When the calibrator cross-correlates exactly $N\log N$ baselines, this fraction falls in the transition region between the two asymptotes. However, assuming that the $N_{\mathrm{bl;ant}}^{-1}$ approximation holds at this fraction of baselines, we can substitute equation (17) into the number of baselines per antenna in equation (14) to get the overall scaling in gain variance:

$$\sigma_g^2 \propto N_{\mathrm{bl;ant}}^{-1} \propto \frac{1}{f\,\frac{N}{2}} \propto \frac{1}{p\log N}. \tag{18}$$

The scaling in gain variance with array size, using subset redundant calibration, is similar to that obtained using low-cadence calibration (equation 13). In both cases, the price one pays for not using $\mathcal{O}(N^2)$ baselines for calibration is that the gain variance scales slower than the case where redundant calibration is performed using the full visibility matrix (equation 7).

Subset redundant calibration leverages the higher constraining power of some baseline-types by allocating computational resources of the calibrator to preferentially cross-correlating those redundant baselines. The least number of baselines that need to be considered is set by the null space of the solution to the redundant-baseline calibration equations. However, using a small fraction of baselines can result in antenna gains that have a non-negligible covariance and location-dependent variances. If the fraction of baselines cross-correlated by a subset redundant calibrator can scale as $\mathcal{O}(N\log N)$ or faster, the gain variance decreases with the increase in array size and scaling in gain variance is similar to that of low-cadence calibration.

6 COMPARISON BETWEEN THE PERFORMANCE OF LOW-CADENCE CALIBRATION AND SUBSET REDUNDANT CALIBRATION

Low-cadence calibration and subset redundant calibration are two potential solutions to calibrating FFT-correlators for redundant array layouts without computing $\mathcal{O}(N^2)$ cross-correlations. In low-cadence calibration, the calibrator computes the visibilities of all baselines through a round-robin of antenna pairs and spends a shorter amount of time on each visibility measurement. As a consequence, the SNR of measured visibilities is lower and leads to a higher variance in the estimated gains. In subset redundant calibration, the calibrator computes the correlations of only a few baselines (preferentially the shorter baseline pairs) and uses these visibilities to estimate the antenna gains. In this case, having fewer measurements leads to higher variance in the estimated gains.

6.1 Scaling in gain variance

Fig. 9 compares the scaling in the average variance of gains estimated using either reduced redundant-baseline calibration method to the scaling in the average variance of gains estimated using redundant-baseline calibration on the full visibility matrix measured at high SNR. For both low-cadence calibration and subset redundant calibration, we assume that the pre-factor $p = 1$ and use only $N\log N$ visibility measurements to estimate antenna gains. The distribution of gain variances, for arrays with number of antennas in the range 1500–5000, is computed using only 16 simulations (instead of 256) due to long simulation times. For the same reason, the distributions for the two largest array sizes with $N > 5000$ antennas are computed using only two simulations each.

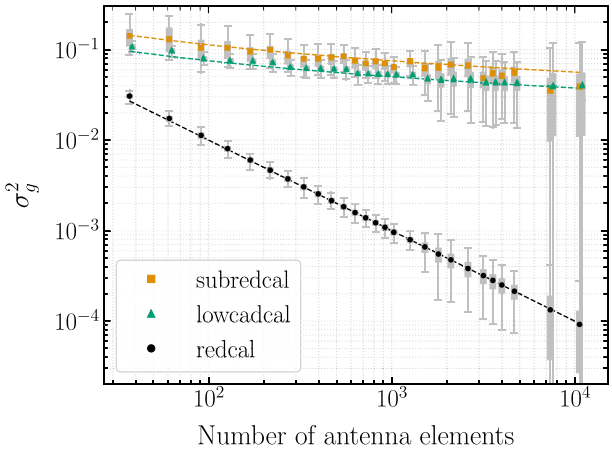

Figure 9. Distributions of antenna gain variance, estimated using low-cadence calibration (`lowcadcal`), subset redundant calibration (`subredcal`) and traditional redundant-baseline calibration (`redcal`) for hexagonal layouts of various sizes. The markers show the mean of the distribution of gain variance for each array size, and the distribution itself is shown using grey solid boxes around each marker (see Fig. 3 for explanation). The larger vertical extent of boxes for arrays with $N > 1500$ elements are due to fewer simulations. The dashed lines represent the theoretically derived trends (equations 7, 13, 18). Low-cadence calibration consistently yields lower gain variance than subset redundant calibration for similar calibrator sizes. While both the reduced redundant-baseline calibration techniques cross-correlate $\sim N\log N$ baselines for each array size, traditional redundant-baseline calibration assumes that the complete visibility matrix with $\sim N^2/2$ baselines is available for calibration. The additional reduction in gain variance for the latter comes at a high computational cost and may not be necessary when scatter in redundant visibilities is dominated by thermal noise (see equations 23 and 24).

For low-cadence calibration, the interval between calibration cycles, t_{cal} is assumed to be constant for all array sizes. The integration time t_{int} has been scaled down according to equation (12) to keep the size of the calibrator at an $\mathcal{O}(N\log N)$ scaling. As predicted, the variance in the estimated antenna gains scales according to equation (13) and is shown by the dashed green line in Fig. 9.

For subset redundant calibration, the SNR of measured visibilities is assumed to be constant for all array sizes. The number of baselines used in the calibration process for a given array size is the closest whole number, to the fraction given by equation (17), that accounts for an integer number of redundant baseline groups in ascending order of their baseline length. That is, baseline-types are considered in the order of their baseline length and added to the subset redundant calibration system only if all the redundant baselines that contribute to that baseline-type can be considered. This causes a non-uniform increase in the number of baseline-types used for calibration as the array size increases. The overall trend in gain variance follows the scaling predicted by equation (18) and is shown by the dashed orange line in Fig. 9.

Subset redundant calibration results in a higher gain variance than low-cadence calibration because, as shown in Fig. 8, the

approximation of inverse measurements per antenna (equation 18) underestimates the gain variance when the pre-factor is unity. However, the scaling predicted by that approximation holds true for all the simulated array sizes. For two largest array sizes in Fig. 9, the estimated gain variance using subset redundant calibration seems comparable to that of low-cadence calibration but is not an exception to this trend. This is because the two simulations used to compute these estimates generated a favourable set of visibilities for subset redundant calibration with the short baselines having larger amplitudes than the long baselines. In general, a subset redundant calibrator would have to cross-correlate more baselines than a low-cadence calibrator to achieve the same gain variance.

Redundant-baseline calibration on the full visibility matrix, measured at high SNR results in a $1/N$ scaling as predicted by equation (7) and shown by the black dashed line in Fig. 9. Such a calibration method requires a calibrator with $\mathcal{O}(N^2)$ computational resources that may not be viable for large arrays. Moreover, the high precision in gains obtained using the full visibility matrix might not be necessary for decreasing the scatter in calibrated redundant visibilities.

6.2 Variance in calibrated redundant visibilities

The visibilities of redundant baselines, that have been calibrated using a reduced redundant-baseline calibration processes, are averaged by the spatial Fourier transform. This converts any post-calibration residual scatter in redundant visibilities into additional noise on the unique visibilities returned by the FFT-correlator.

If the estimated antenna gains are exactly the true gains, the scatter in redundant-baseline averaged visibilities is given by σ^2/N_α, assuming that the variance in the thermal noise of visibilities is similar for all baselines and represented by σ^2. However, gains estimated using redundant-baseline calibration diverge from the true gains with an average scatter that is given by equation (6). Hence, the calibrated redundant visibilities have a residual scatter that comes from both the thermal noise in the measurements and the variance in the estimated gains.

The variance in calibrated visibilities can be derived using equation (2) and the first-order approximation for the variance of non-linear functions. In this derivation, we have assumed that the calibrated visibilities are not correlated with each other. Moreover, since the gains estimated using reduced redundant-baseline calibration are applied to a different set of visibilities that those used to estimate them, we can also ignore the covariance between visibilities and gains. The multiplying antenna gains, however, have a non-negligible covariance that is represented by the terms $\rho_{g_i g_j}$ in the following equation. This represents the off-diagonal components in the covariance matrix $\mathbf{C}'$ in equation (16):

$$\sigma^2_{V_\alpha^{\rm unique}} \approx \frac{\left|V_\alpha^{\rm unique}\right|^2}{N_\alpha^2} \sum_{(i,j)\in\alpha} \left[\frac{\sigma^2_{ij}}{|V_{ij}|^2} + \frac{\sigma^2_{g_i}}{|g_i|^2} + \frac{\sigma^2_{g_j}}{|g_j|^2} + 2\frac{\rho_{g_i g_j}}{|g_i g_j|} \right]. \tag{19}$$

We can simplify this further under the assumptions that the average amplitude of all antenna gains is close to one ($|g_i|^2 \sim 1$) and that the gain variance of all antennas is similar and given by σ^2_g. Note that the relative variance of visibilities is just the inverse squared SNR. In the equation below, $\mathrm{SNR}_{\alpha;\,\mathrm{unique}}$ is the SNR of the visibility computed by the FFT-correlator for the baseline-type represented by α and $\mathrm{SNR}_{\mathrm{full}}$ is the average SNR of the N_α number of redundant visibilities that belong to that baseline-type:

$$(\mathrm{SNR})^{-2}_{\alpha;\,\mathrm{unique}} \approx \frac{1}{\mathrm{N}_\alpha} \left[(\mathrm{SNR})^{-2}_{\mathrm{full}} + 2\sigma^2_{\mathrm{g}} + \frac{2}{\mathrm{N}_\alpha} \sum_{(\mathrm{i,j})\in\alpha} \rho_{\mathrm{g_i g_j}} \right]. \tag{20}$$

This equation gives the acceptable range of gain variance and covariance for the gains estimated using redundant-baseline calibration. When the gain variance and covariance is much smaller than the thermal noise in visibilities, the first term dominates the residual scatter. If this is satisfied, lowering gain variance by using a larger calibrator will not improve the variance in calibrated visibilities.

6.2.1 *Variance in redundant visibilities calibrated using low-cadence calibration*

The variance in antenna gains estimated using low-cadence calibration depends on the SNR of the visibilities computed in the calibrator which in turn depends on the integration time (equation 11) available for each cycle of computation. The relationship between the SNR of the full visibility matrix and that of the reduced visibility matrix computed in by calibrator, that scales as $\mathcal{O}(N \log N)$, can be written using equation (12) as

$$\frac{(\mathrm{SNR})^{-2}_{\mathrm{full}}}{(\mathrm{SNR})^{-2}_{\mathrm{reduced}}} = \frac{t_{\mathrm{int;reduced}}}{t_{\mathrm{int;full}}} = \left(\frac{2p\log N}{N}\right) \frac{t_{\mathrm{cal}}}{t_{\mathrm{int;full}}}, \tag{21}$$

where $t_{\mathrm{int;\,full}}$ is the integration time used in an FX- or FFT-correlator. Substituting the above relation into equation (6), we can write the variance in estimated gains in terms of the SNR of the visibility matrix that the gains calibrate:

$$\begin{aligned} \sigma^2_g &= (\mathrm{SNR})^{-2}_{\mathrm{reduced}} \left(\mathbf{A}^\dagger\mathbf{A}\right)^{-1}_{(N\times N)} \\ &= (\mathrm{SNR})^{-2}_{\mathrm{full}} \left(\frac{N}{2p\log N}\right) \frac{t_{\mathrm{int;full}}}{t_{\mathrm{cal}}} \frac{1}{N/2} \\ &= (\mathrm{SNR})^{-2}_{\mathrm{full}} \left(\frac{1}{p\log N}\right) \frac{t_{\mathrm{int;full}}}{t_{\mathrm{cal}}}. \end{aligned} \tag{22}$$

When t_{cal} is large, the SNR of the reduced visibility matrix is larger and the corresponding gain variance is smaller. Assuming the case where the interval of calibration is same as the integration time in the FFT correlator and substituting the above equation into equation (20) we get the following for the variance in redundant visibilities:

$$(\mathrm{SNR})^{-2}_{\mathrm{unique}} \approx \frac{(\mathrm{SNR})^{-2}_{\mathrm{full}}}{\mathrm{N}_\alpha} \left[1 + \left(\frac{2}{\mathrm{p}\log \mathrm{N}}\right)\right]. \tag{23}$$

When redundant-baseline calibration is performed with the full visibility matrix, the gain covariance terms in equation (20) are around an order of magnitude smaller than the variance term so we drop the third term for clarity. In large arrays, where an FFT-correlator architecture would be preferable to an FX-correlator, the contribution of gain variance (second term) to the variance in calibrated redundant visibilities is smaller than the thermal noise in the measure visibilities (first term). Hence, the precision in gain variance obtained from using an $\mathcal{O}(N \log N)$ calibrator is sufficient for the purpose of calibrating voltages for an FFT-correlator.

6.2.2 *Variance in redundant visibilities calibrated using subset redundant calibration*

In subset redundant calibration, the integration time does not change between the calibrator and the FFT-correlator. Consequently, the SNR of measured visibilities is the same for both data sets. However,

the number of baselines used in the calibration process is lower, resulting in a higher variance in the estimated gains. For a calibrator that scales as $\mathcal{O}(N \log N)$ with array size, the relationship between gain variance and number of baselines is given by equation (18) where the proportionality constant is the SNR of visibilities. Substituting equation (18) into equation (20), we get

$$(\mathrm{SNR})^{-2}_{\mathrm{unique}} \approx \frac{(\mathrm{SNR})^{-2}_{\mathrm{full}}}{\mathrm{N}_\alpha}\left[1 + \frac{2}{\mathrm{p}\log \mathrm{N}} + \frac{2}{\mathrm{N}_\alpha}\sum_{(\mathrm{i,j})\in\alpha}\rho_{\mathrm{g_i g_j}}\right]. \quad (24)$$

The gain covariance terms $\rho_{g_i g_j}$, for gains estimated using subset redundant calibration, are sometimes comparable to the gain variance terms. When the pre-factor (p) is small, the amplitude of the covariance scales similarly to the gain variance with increase in size of the array. This effectively doubles the contribution of the variance term in the above equation. When the pre-factor is large, the covariance is only a small fraction of the gain variance and can be ignored. Overall, the contribution of the second and third terms in the above equation is much smaller than the thermal noise in measured visibilities for any reasonably large array. Hence, the precision in gains estimated using subset redundant calibration is also sufficient for the purpose of calibrating voltages for an FFT-correlator.

Fig. 9 might give a misleading impression that traditional redundant-baseline calibration is superior to either of the reduced redundant-baseline schemes, by providing gains that have orders-of-magnitude lower variance. However, this additional gain precision comes at a high computational cost and might not be necessary for large arrays where the contribution of gain variance to the overall scatter in redundant visibilities is only a small fraction of the thermal noise.

6.3 Bias in estimated variables

The gains estimated using either reduced redundant-baseline calibration process are unbiased estimates of the true value. This can be verified through simulations that have constant underlying gains and visibilities, and different realizations of the noise in the measured visibilities. Averaging the solutions obtained over multiple such simulations decreases the noise in the estimated parameters and can expose an underlying bias, if any. Fig. 10 shows the deviation in averaged gains from the input true gains, normalized by the variance expected in the gains. The errorbars represent the antenna-to-antenna variation which also averages down. Gains that have been averaged over N_{sim} independent noise realizations have a factor of $1/N_{\mathrm{sim}}$ smaller deviation, which is expected when the estimated gains differ from truth only within the Gaussian random noise in the measurement. This trend is marked by the dashed black line in the figure. A bias in gains, within the precision exposed by averaging down noise, would have resulted in a deviation from this trend.

The visibilities computed by the FFT-correlator $\boldsymbol{V}^{\mathrm{unique}}$ are also unbiased when the visibility matrix used for reduced redundant-baseline has a high SNR. The deviation of the calibrated and redundant-baseline averaged visibilities from the simulated input, averages down according to the trend expected for Gaussian random noise. However, empirically, we find that the calibrated visibilities are sometimes biased when the variance in estimated gains is larger than $\sim 10^{-5}$. It is possible that there is a low level of bias, that is not exposed by averaging over $N_{\mathrm{sim}} = 4096$ simulations, in the visibilities calibrated with lower variance gains. To clarify, the gain variance plotted in all the figures in this paper has been normalized by the thermal noise in the visibilities used for calibration. The absolute

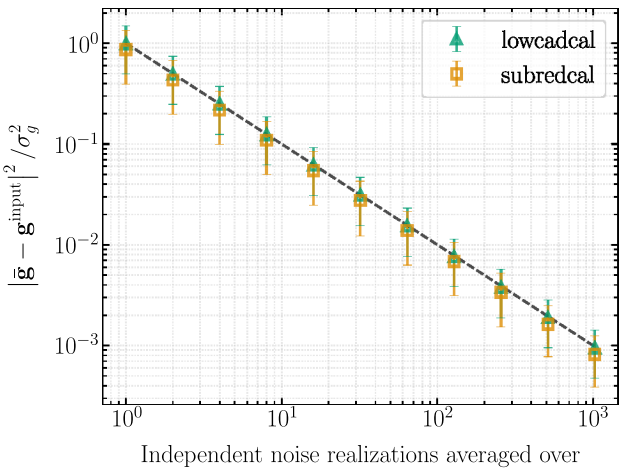

Figure 10. Deviation of averaged gains from their true value for low-cadence calibration (`lowcadcal`; green triangles) and subset redundant calibration (`subredcal`; orange squares). The markers are the result of a simulation with the same underlying input variables (both gains and unique visibilities) but different noise realizations of the measured visibilities. The x-axis shows the number of such simulations that the variables have been averaged over. The errorbars reflect the antenna-to-antenna variation. The dashed line represents the trend expected when the estimated gains differ from truth only within Gaussian noise. Both low-cadence calibration and subset redundant calibration yield unbiased gain solutions.

gain variance in all these simulations is less than the empirical value above which we find a bias in the calibrated visibilities.

Assuming this threshold in gain variance is real, it translates to a requirement of an SNR $\gtrsim 20$ in the measured visibilities for a 300 antenna array when implementing low-cadence calibration, and higher when implementing subset redundant calibration. This minimum SNR requirement decreases as $1/\sqrt{N}$ with increase in array size and may not be an issue for large-N arrays.

This requirement of a minimum SNR in the reduced visibility matrix could be due two possible reasons: (a) if the assumption that the non-linear equations of redundant-baseline calibration can be optimized by solving the linearized system of equations does not hold at this limit or (b) if the product of gains form an asymmetric distribution about their mean value and do not average down. If the former is true, a linearized solver of equation (1) would result in a biased solution when the noise in measured visibilities is high. However, we find that linearization based on Taylor expansion of variables and `omnical` always result in unbiased gains and visibilities irrespective of the SNR in the measured visibilities.

The more favourable explanation seems to be latter. As shown by O'Donoughue & Moura (2012), the probability distribution function (PDF) of the product of two complex Gaussian random variables is not a simple Gaussian distribution. Moreover, the resulting distribution can be asymmetric if two complex random variables are drawn from a non-zero mean Gaussian, as is the case with antenna gains.

6.4 χ^2_r of estimated gains and visibilities

As was discussed in Section 2.2, the gains computed using a reduced redundant-baseline calibration scheme are estimated from the visibility matrix $\mathbf{V}^{\mathrm{reduced}}$ but applied to a different visibility matrix $\mathbf{V}^{\mathrm{unique}}$. The redundant-baseline calibration process is designed to minimize the χ^2_r between the visibility matrix used for calibration and the estimated variables. In the case of reduced redundant-baseline

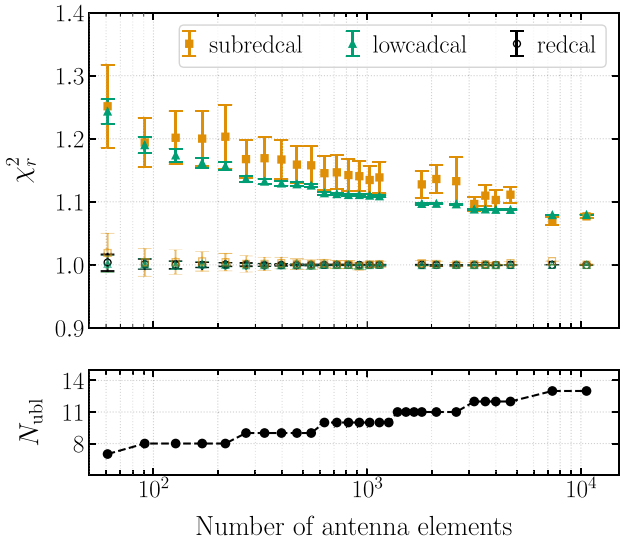

Figure 11. The upper panel shows the goodness of fit of the gains and unique visibilities that have been estimated using either reduced redundant-baseline calibration scheme. The hollow points (orange and green) show the fit of the estimated gains to the visibilities computed by the calibrator $\mathbf{V}^{\rm reduced}$ and is close to one as expected. The hollow black points show the fit of gains estimated using full redundant-baseline calibration to $\mathbf{V}^{\rm full}$. The solid orange and green points show the χ_r^2 of the fit to the full visibility matrix $\mathbf{V}^{\rm full}$ and these are also not far from one. The associated errorbars show the range of χ_r^2 obtained over multiple simulations. The jump in χ_r^2 for low-cadence calibration arises from a non-uniform increase in the calibrator size. The number of unique baseline-types that are cross-correlated by the calibrator are shown in the lower panel as a proxy for the size of the calibrator.

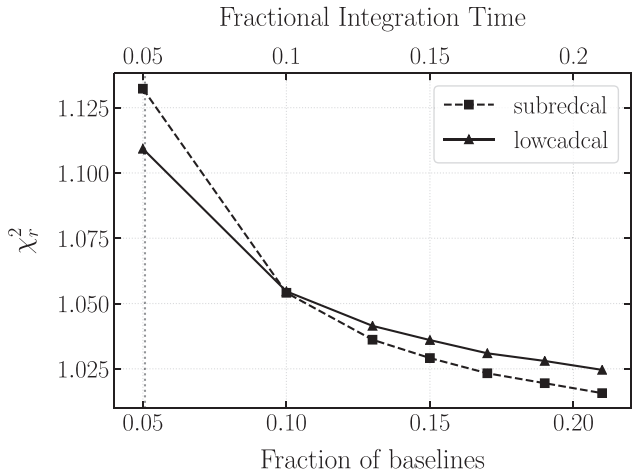

Figure 12. Reduced chi-squared as a function of the computational resources allocated to a reduced redundant calibrator for a 331 antenna array. The size of the calibrator is marked in the fraction of baselines that can be processed for subset redundant calibration (bottom axis) or the fraction of integration time that can be spent on each baseline for low-cadence calibration (top axis). The vertical dotted line marks the size of a $\mathcal{O}(N \log N)$ calibrator. At larger calibrator sizes, subset redundant calibration yields better gain estimates than low-cadence calibration.

calibration, this is the χ_r^2 evaluated between the estimated gains, model visibilities that are discarded by the calibrator and the reduced visibility matrix. In Fig. 11, this χ_r^2 is represented by hollow orange and green markers for subset redundant calibration and low-cadence calibration, respectively, and has the expected value of one.

The χ_r^2 estimated using gains computed by the reduced redundant-baseline calibration process, unique visibilities computed by the FFT-correlator and the full visibility matrix is a better metric to assess the effectiveness of the calibration process. This χ_r^2 is represented by the solid orange and green points in Fig. 11 for subset redundant calibration and low-cadence calibration respectively. The χ_r^2 for array sizes with number of antennas in the range of 1500–5000 is estimated over only 16 simulations, rather than 256, due to long simulation times. For the same reason, the χ_r^2 for the two largest array sizes with $N > 5000$ antennas is estimated using only two simulations.

Both the reduced redundant-baseline calibration schemes yield a χ_r^2 that is close to one, indicating that the estimated parameters are a reasonable fit to the full visibility matrix. The larger χ_r^2 of subset redundant calibration as compared to low-cadence calibration, and the larger range of χ_r^2 obtained over multiple simulations (represented by the errorbars) is a direct consequence of the higher gain variance in the former compared to the latter. The non-smooth trend in the χ_r^2 of low-cadence calibration is due to a non-uniform increase in the size of the calibrator used in simulation. The jump in χ_r^2 of low-cadence calibration is correlated with the increase in number of unique-baseline-types processed by the calibrator, because this leads to a jump in the integration time available to each cycle of correlation in the calibrator.

Fig. 12 compares the performance of low-cadence calibration to subset redundant calibration for calibrators of various sizes operating a fixed array size. When the resources allocated to the calibrator are sufficient to cross-correlate exactly $N \log N$ baselines, which was the assumption throughout this paper, low-cadence calibration yields better gain estimates. This calibrator size is marked by a vertical dotted line in the figure. However, at larger calibrator sizes, subset redundant calibration performs better than low-cadence calibration. This is because the subset redundant calibrator preferentially spends time on baselines that have a higher constraining power and integrates them to a higher SNR than a low-cadence calibrator of a similar size.

7 CONCLUSION

Low-frequency radio interferometers with $N \gg 1000$ antennas are being proposed for targeting the 21 cm signal at cosmological distances. These arrays will face a significant computational cost in building traditional FX-correlators that scale as $\mathcal{O}(N^2)$ with the number of antennas in the array. FFT-correlators attempt to decrease this scaling to $\mathcal{O}(N \log N)$, thereby decreasing the cost of the correlator backend. However, FFT-correlators only produce meaningful output when the input antenna voltages are calibrated. In the past, the few experiments that have used FFT-correlators on redundant arrays had an FX-correlator working in parallel for computing the visibilities required for calibration. This is, however, a non-scalable solution for large-N arrays. In this paper, we propose a $\mathcal{O}(N \log N)$ calibrator design that can operate in parallel with the FFT-correlator, forming a self-contained correlator system that can be scaled to large-N arrays built on a regular grid. We discuss two calibration algorithms, that can be employed by such a calibrator, which are modifications of redundant-baseline calibration.

Redundant-baseline calibration relies on the array having multiple redundant baselines or antenna pairs with the same displacement vector. Non-redundancies in the array, originating from beam variation, antenna placement errors, etc., could lead to errors in the gain solution as discussed by Orosz et al. (2018) and Byrne et al. (2019). These gain errors are likely to affect low-cadence calibration and subset redundant calibration, even though the FFT-correlator

averages theoretically redundant baselines together. An analysis of the impact of non-redundancy on an FFT-correlated telescope is deferred for future work.

In low-cadence calibration, antenna pairs are cross-correlated in a round-robin fashion over multiple integration cycles, producing the full visibility matrix once every given number of cycles. However, the time period between two calibration cycles $t_{\rm cal}$ is fixed by the inherent gain variability of the array. This leads to a decrease in the integration time available for each cycle, as the array size increases. A decrease in integration time results in lower SNR visibilities within the calibrator, and consequently a higher variance in the estimated gains. For a calibrator that is restricted to an $\mathcal{O}(N \log N)$ scaling, the gain variance scales as $1/\log N$ with increasing array size. The contribution of this gain variance to the overall variance in the visibilities computed by the FFT-correlator is much smaller than the contribution of thermal noise in measurements. Hence, low-cadence calibration is a suitable calibration scheme to estimate gains that are effective in calibrating voltages for the FFT-correlator.

In subset redundant calibration, the calibrator computes the visibilities of only a few baseline groups without compromising on their SNR. By nature of the system of equations solved by redundant-baseline calibration, the larger baseline groups contribute higher constraining power to antenna gains. Subset redundant calibration exploits this property and computes cross-correlations of only these baseline groups. The minimum number of baseline-types that need to be considered for subset redundant calibration is determined by the null space of the solution set, which should ideally not have more than the four known degeneracies. However, using only a small fraction of all the baselines, which satisfy this degeneracy criterion, could still result in antenna gains that have a high variance and covariance. A calibrator that can cross-correlate $N \log N$ baselines results in antenna gains that have a $1/\log N$ scaling in variance as well. As in the case of low-cadence calibration, the variance of the estimated antenna gains forms only a small fraction of the total scatter in visibilities computed by the FFT-correlator. Hence, subset redundant calibration is also a suitable scheme of calibration for estimating gains that can minimize scatter in redundant visibilities.

The gains, estimated using either reduced redundant-baseline calibration method, are unbiased and converge to the true value when noise averaged. When comparing low-cadence calibration and subset redundant calibration, we find that low-cadence calibration consistently yields lower variance antenna gains when the size of the calibrator is held constant. This consequently leads to a better fit to measured visibilities and the χ_r^2 estimated using the variables computed by low-cadence calibration is lower than that of subset redundant calibration. However, for large-N arrays low-cadence calibration could involve optimizing over a few hundred thousand equations which could potentially hamper the real-time nature of the calibration parameters required. If the calibrator for a given array has more computational resources, subset redundant calibration can result in better antenna gain estimates. Ultimately, the calibration method that is suitable for a large array depends on the antenna design parameters, the array layout, computational resources available and the science goal at hand.

ACKNOWLEDGEMENTS

We would like to thank Jack Hickish, Phil Bull, Adrian Liu, Nithyanandan Thyagarajan, Adam Beardsley, and Judd Bowman for valuable discussions that contributed to this paper. We would also like to thank the anonymous reviewer for providing suggestions that improved the readability and value of this paper. This material is based upon work supported by the National Science Foundation under grants #1636646 and #1836019. This research is funded in part by the Gordon and Betty Moore Foundation. JSD gratefully acknowledges the support of the NSF AAPF award #1701536 and the Berkeley Center for Cosmological Physics.

DATA AVAILABILITY

The data underlying this article has been generated using code publicly available in the repository: https://github.com/HERA-Team/hera_sim.

REFERENCES

Ali Z. S. et al., 2015, ApJ, 809, 61
Bandura K. et al., 2014, SPIE, 9145 914522
Beardsley A. P., Thyagarajan N., Bowman J. D., Morales M. F., 2017, MNRAS, 470, 4720
Byrne R. et al., 2019, ApJ, 875, 70
Cooley J. W., Tukey J. W., 1965, Math. Comput., 19, 297
Daishido T. et al., 1991, ASPCS, 86
DeBoer D. R. et al., 2017, PASP, 129, 045001
Dillon J. S., Parsons A. R., 2016, ApJ, 826, 181
Dillon J. S. et al., 2018, MNRAS, 477, 5670
Dillon J. S. et al., 2020 MNRAS, submitted
Ellingson S. W., Kassim N., Perez M., Junor W., Henning P., 2005, ASPCS, 345, 321
Escoffier R. P. et al., 2007, A&A, 462, 801
Foster G., Hickish J., Magro A., Price D., Zarb Adami K., 2014, MNRAS, 439, 3180
Kent J., Dowell J., Beardsley A., Thyagarajan N., Taylor G., Bowman J., 2019, MNRAS, 486, 5052
Kern N. S., Parsons A. R., Dillon J. S., Lanman A. E., Fagnoni N., de Lera Acedo E., 2019, ApJ, 884, 105
Kern N. S. et al., 2020, ApJ, 890, 122
Kolopanis M. et al., 2019, ApJ, 883, 133
Li W. et al., 2018, ApJ, 863, 170
Liu A., Shaw J. R., 2020, PASP, 132, 062001
Liu A., Tegmark M., Morrison S., Lutomirski A., Zaldarriaga M., 2010, MNRAS, 408, 1029
Marthi V. R., Chengalur J., 2014, MNRAS, 437, 524
Morales M. F., 2011, PASP, 123, 1265
Newburgh L. B. et al., 2014, SPIE, 9145, 91454V
Noorishad P., Wijnholds S. J., van Ardenne A., van der Hulst J. M., 2012, A&A, 545, A108
O'Donoughue N., Moura J. M. F., 2012, IEEE Trans. Signal Proc., 60, 1050
Orosz N., Dillon J. S., Ewall-Wice A., Parsons A. R., Thyagarajan N., 2019, MNRAS, 487 537
Parsons A. R. et al., 2010, AJ, 139, 1468
Tegmark M., Zaldarriaga M., 2009, Phys. Rev. D, 79, 083530
Tegmark M., Zaldarriaga M., 2010, Phys. Rev. D, 82, 103501
Thompson A. R., Clark B. G., Wade C. M., Napier P. J., 1980, ApJS, 44, 151
Thompson A. R., Moran J. M., Swenson George W. J., 2017, Interferometry and Synthesis in Radio Astronomy, 3rd edn., Springer, Cham
Thyagarajan N., Beardsley A. P., Bowman J. D., Morales M. F., 2017, MNRAS, 467, 715
Tingay S. J. et al., 2013, PASA, 30, e007
van Haarlem M. P. et al., 2013, A&A, 556, A2
Wieringa M. H., 1992, Exp. Astron., 2, 203
Zheng H. et al., 2014, MNRAS, 445, 1084
Zheng H. et al., 2017, MNRAS, 465, 2901

APPENDIX: HIERARCHICAL CALIBRATION

Hierarchical redundant-baseline calibration was originally presented by Zheng et al. (2014, appendix B). It is based on separating a

redundant array into sub-arrays and cross-correlating all the antenna pairs in each sub-array. Antenna gains are estimated by performing redundant-baseline calibration on each sub-array independently. The degenerate parameters of each sub-array are tied together by choosing one antenna from each sub-array and redundantly calibrating the array they form. For redundant arrays where the correlator also operates hierarchically (Tegmark & Zaldarriaga 2010), such a calibrator system would parallel the correlator layout in the field and reduce networking. However, this is not an optimal calibration solution for FFT-correlators that operate on the entire array.

The spatial Fourier transform in an FFT-correlator can only be performed on a machine containing the voltages of all antennas in the array. This requires a large corner-turn as in FX-correlators, where the voltages from all antennas over a narrow bandwidth are collected at a central server using an Ethernet switch or a similar device. When data from all the antennas are available at a central location, the baselines that yield optimal gain solutions can be chosen in any fashion for the purpose of calibration. While hierarchical calibration has the potential to simplify the networking required for large arrays, the networking required for FFT-correlators renders this simplification moot. For this reason, hierarchical calibration has been classified as a special case of subset redundant calibration and is not presented in more detail.

This paper has been typeset from a TEX/LATEX file prepared by the author.