

Baseline Dependent Averaging in Radio Interferometry

S. J. Wijnholds^{1*}, A. G. Willis², S. Salvini³

¹*ASTRON, Oude Hoogeveensedijk 4, 7991 PD Dwingeloo, The Netherlands*

²*National Research Council of Canada, P.O. Box 248, Penticton, BC V2A 6J9, Canada*

³*OeRC, 7 Keble Road, Oxford, OX1 3QG, UK*

26 February 2018

ABSTRACT

This paper presents a detailed analysis of the applicability and benefits of baseline dependent averaging (BDA) in modern radio interferometers and in particular the Square Kilometre Array (SKA). We demonstrate that BDA does not affect the information content of the data other than a well-defined decorrelation loss for which closed form expressions are readily available. We verify these theoretical findings using simulations. We therefore conclude that BDA can be used reliably in modern radio interferometry allowing a reduction of visibility data volume (and hence processing costs for handling visibility data) by more than 80%.

Key words: Instrumentation: interferometers, Methods: analytical, Methods: numerical, Techniques: interferometric

1 INTRODUCTION

Managing the large amount of visibility data is one of the biggest challenges in imaging science with modern radio interferometers, such as the Karl G. Jansky Very Large Array (JVLA, Perley et al. (2009)), the Murchison Widefield Array (MWA, Tingay et al. (2013)), the Low Frequency Array (LOFAR, van Haarlem et al. (2013)) and in particular the future Square Kilometre Array (SKA, Dewdney et al. (2009)). The visibility data volumes are driven upwards by the small channel bandwidth and short correlator dump times required to limit smearing and decorrelation effects on the longest baselines and at the edges of the field-of-view to an acceptable level (Bridle & Schwab 1999). Since this small channel width and high time resolution are not required on the shorter baselines, Cotton (2009) and Skipper (2014) have considered baseline dependent averaging (BDA) in the context of the VLA and the SKA respectively. BDA has also been used for MWA, which initially had to operate with stringent storage capacity restrictions (Mitchell et al. 2008). Besides data compression, BDA can also be used to shape the field-of-interest (Atemkeng et al. 2016).

Time and frequency averaging causes smearing effects that do not have a convolutional nature in the image plane as explained by Bridle & Schwab (1999). Cotton (2009) shows that this has detrimental effects on the achievable dynamic range, which is of particular concern in the context of modern radio interferometers. He alludes to the possibility to make first order corrections for these smearing effects at the expense of an increase in computing requirements. Since

LOFAR experience shows that imaging pipelines for SKA may actually be limited by data handling instead of computing, this increase in required computing may be acceptable if it makes the visibility data volumes more manageable. The question then still remains, whether BDA poses a fundamental limit on dynamic range if the appropriate corrections can be made during the imaging process. This is the first question considered in this paper. It will be addressed by an analysis of information content of visibilities with and without application of BDA. This analysis is presented in Sec. 2.

The length of the solution interval for calibration is another concern in the context of BDA, in particular for time averaging. In this case, the integration time on the shortest baselines may become longer than the length of the solution interval required to keep track of the fastest changes in the observing system. The obvious solution would be to limit the maximum allowed integration time on any baseline to an interval that is shorter than the solution interval used for calibration. The problem with this approach is, that the length of the solution interval may be chosen differently during post-processing of the data to deal with, for example, large temperature variations or worse ionospheric or tropospheric conditions than expected. This risk can be mitigated by the Compress-Expand-Compress (CEC) method described by Salvini & Wijnholds (2017). For convenience of the reader, we recapitulate the main results presented in Salvini & Wijnholds (2017) in Sec. 3. Based on these results, we will assume that we are justified in averaging the data to the extent that we do in our analysis and simulations.

We validate our analysis using simulations in which we study the difference between images obtained from full-

* E-mail: wijnholds@astron.nl

resolution visibilities and visibilities obtained after applying BDA. These simulations, which are presented in Sec. 4, are set up in such a way that the impact of BDA is isolated from other effects, such as incomplete sky models or imperfect instrument models, that we usually face when making high dynamic range images. These simulations confirm the main conclusion from our theoretical analysis that BDA is not likely to have a detrimental effect on imaging performance besides a well-understood decorrelation effect.

2 THEORETICAL ANALYSIS

2.1 Mathematical description of BDA

A correlator usually correlates narrowband time-series data from P receiving elements, which can either be aperture array stations or dishes. We can write the signal from the p th receiving element at frequency ν as $x_p((LN+n)T, \nu)$, where T is the sample interval, n is the sample index in the l th short term integration (STI) and N is the length of a short term integration, which is determined by the correlator dump time needed to accurately sample the visibilities on the longest baselines. For Nyquist sampled input signals, $N = B\tau$, where B is the channel bandwidth and τ is the integration time of a STI. The corresponding time samples for all P receiving elements can be stacked into the array signal vector

$$\mathbf{x}((LN+n)T, \nu) = \begin{bmatrix} x_1((LN+n)T, \nu) \\ \vdots \\ x_P((LN+n)T, \nu) \end{bmatrix} \in \mathbb{C}^{P \times 1}. \quad (1)$$

The time samples for the l th STI can be collected in the matrix

$$\mathbf{X}_{l,\nu} = [\mathbf{x}_\nu((LN+1)T), \dots, \mathbf{x}_\nu((LN+N)T)] \in \mathbb{C}^{P \times N}, \quad (2)$$

where we have changed the frequency ν into an indexing parameter, since BDA can not only be applied along the time axis, but also along the frequency axis. Using these samples, a correlator that does not use BDA, will estimate the visibility matrix or array covariance matrix

$$\hat{\mathbf{R}}_{l,\nu} = \frac{1}{N} \mathbf{X}_{l,\nu} \mathbf{X}_{l,\nu}^H \in \mathbb{C}^{P \times P} \quad (3)$$

for the l th STI. Note that the autocorrelations are included in this matrix. Although these are often discarded in radio astronomical data reduction, they are essential for our statistical analysis, because the autocorrelations measure the system temperature.

Each of these visibility matrices can be vectorized by stacking its columns to form visibility vectors $\hat{\mathbf{r}}_{l,\nu} = \text{vec}(\hat{\mathbf{R}}_{l,\nu}) \in \mathbb{C}^{P^2 \times 1}$. We can replace the indexing parameters l and ν to index STIs over time and frequency respectively by a single index parameter k . Let us assume that K STIs, which can be collected over time or frequency or both, are combined in a self-calibration and imaging process to estimate calibration and image parameters stacked in the parameter vector $\boldsymbol{\theta}$. After defining the raw visibility data vector as

$$\hat{\mathbf{r}} = [\hat{\mathbf{r}}_1^T, \dots, \hat{\mathbf{r}}_K^T] \in \mathbb{C}^{KP^2 \times 1}, \quad (4)$$

the self-calibration problem can be formulated as

$$\hat{\boldsymbol{\theta}} = \underset{\boldsymbol{\theta}}{\text{argmin}} \|\hat{\mathbf{r}} - \mathbf{r}(\boldsymbol{\theta})\|^2, \quad (5)$$

where $\mathbf{r}(\boldsymbol{\theta})$ represents the parameterized measurement equation.

The length of the raw visibility data vector can be reduced from KP^2 to $M \leq KP^2$ by averaging the raw visibilities on some baselines. Integration over multiple STIs by *summation* can be described by

$$\mathbf{r}_{\text{sum}} = \mathbf{I}_s \mathbf{r} \in \mathbb{C}^{M \times 1}. \quad (6)$$

The selection matrix $\mathbf{I}_s \in \mathbb{R}^{M \times KP^2}$ is a sparse matrix with only a single element equal to unity in each column. Each row contains at least a single element equal to unity, but may contain multiple elements with value 1 depending on the number of raw visibility samples being added into a single summed visibility. This selection matrix can be partitioned in selection matrices for each individual STI $\mathbf{I}_{s,k} \in \mathbb{R}^{M \times P^2}$ as

$$\mathbf{I}_s = [\mathbf{I}_{s,1}, \dots, \mathbf{I}_{s,K}]. \quad (7)$$

The selection matrix for a single STI may have one or more rows containing only zeros signifying that that particular STI does not contribute to the summed visibility associated with that row.

The *averaged* visibilities \mathbf{r}_{ave} follow from the summed visibilities by dividing each element of \mathbf{r}_{sum} by the number of raw visibilities that contributed to that element. This can be done using the diagonal weighting matrix

$$\mathbf{W} = (\mathbf{I}_s \mathbf{I}_s^T)^{-1} \in \mathbb{R}^{M \times M} \quad (8)$$

to obtain

$$\mathbf{r}_{\text{ave}} = \mathbf{W} \mathbf{r}_{\text{sum}} = \mathbf{W} \mathbf{I}_s \mathbf{r} \in \mathbb{C}^{M \times 1}. \quad (9)$$

Ignoring decorrelation effects, which will be discussed in Sec. 2.4, the visibility values that are averaged together, are assumed to measure the same physical quantity, i.e., the expected value of these data values is assumed to be the same. This assumption means that the expected value of \mathbf{r} can be obtained from the expected value of $\hat{\mathbf{r}}_{\text{ave}}$ by

$$\mathbf{r} = \mathcal{E}\{\hat{\mathbf{r}}\} = \mathbf{I}_s^H \mathbf{r}_{\text{ave}}. \quad (10)$$

Strictly speaking, this equation is not valid for realized values, because the noise on the non-averaged visibilities is averaged in the averaged visibilities and can therefore not be recovered. However, as long as we work with expected values, we can use this equation to transform from the space of averaged visibilities to the space of non-averaged visibilities. Interestingly, (10) describes the most basic expand step in the Compress-Expand-Compress approach proposed by Salvini & Wijnholds (2017) and summarized in Sec. 3.

2.2 Information content of averaged visibilities

The Cramer-Rao bound (CRB) provides a lower bound on the covariance of all estimated parameters for an unbiased estimator (Kay 1993). The CRB is the inverse of the Fisher Information Matrix (FIM), which measures the amount of information a given data vector contains about an unknown parameter vector. In this section, we assess the preservation

of information when applying BDA by comparing the CRB for estimation of \mathbf{r}_{ave} from \mathbf{r} with the covariance of \mathbf{r}_{ave} obtained by propagating the noise on the raw visibilities \mathbf{r} to the averaged visibilities.

The FIM is defined as (Kay 1993)

$$\mathbf{F} = \mathbf{J}^H \text{cov}(\mathbf{r})^{-1} \mathbf{J}, \quad (11)$$

where $\mathbf{J} = \partial \mathbf{r} / \partial \boldsymbol{\theta}^T$ is a Jacobian matrix containing the partial derivatives for each element of \mathbf{r} to all parameters in the parameter vector $\boldsymbol{\theta}$ and $\text{cov}(\mathbf{r})$ is the covariance matrix of the raw visibility vector \mathbf{r} . For our analysis, we have $\boldsymbol{\theta} = \mathbf{r}_{\text{ave}}$, so, using (10), the Jacobian matrix becomes

$$\mathbf{J} = \frac{\partial \mathbf{r}}{\partial \mathbf{r}_{\text{ave}}^T} = \mathbf{I}_s^H. \quad (12)$$

The CRB for estimating \mathbf{r}_{ave} from \mathbf{r} thus becomes

$$\mathbf{C} = \mathbf{F}^{-1} = \left(\mathbf{I}_s \text{cov}(\mathbf{r})^{-1} \mathbf{I}_s^H \right)^{-1}. \quad (13)$$

If the STIs are uncorrelated, i.e., if the time and frequency samples can be considered as independent measurements, the covariance of \mathbf{r} can be written as

$$\text{cov}(\mathbf{r}) = \begin{bmatrix} \text{cov}(\mathbf{r}_1) & & \\ & \ddots & \\ & & \text{cov}(\mathbf{r}_K) \end{bmatrix}, \quad (14)$$

i.e., as a block diagonal matrix with the covariance matrices of the individual STIs on the main diagonal. Using the fact that the inverse of a block diagonal matrix is a block diagonal matrix formed by the inverse of the individual blocks and the fact that \mathbf{I}_s can be partitioned per STI as shown in (7), we can write (13) as

$$\mathbf{C} = \left(\sum_{k=1}^K \mathbf{I}_{s,k} \text{cov}(\mathbf{r}_k)^{-1} \mathbf{I}_{s,k}^H \right)^{-1}. \quad (15)$$

Using standard error propagation formulas, we find that the covariance of \mathbf{r}_{ave} obtained using the BDA procedure described earlier is given by

$$\begin{aligned} \text{cov}(\mathbf{r}_{\text{ave}}) &= \left(\frac{\partial \mathbf{r}_{\text{ave}}}{\partial \mathbf{r}^T} \right) \text{cov}(\mathbf{r}) \left(\frac{\partial \mathbf{r}_{\text{ave}}}{\partial \mathbf{r}^T} \right)^H \\ &= \mathbf{W} \mathbf{I}_s \text{cov}(\mathbf{r}) \mathbf{I}_s^H \mathbf{W}^H \\ &= \mathbf{W} \left(\sum_{k=1}^K \mathbf{I}_{s,k} \text{cov}(\mathbf{r}_k) \mathbf{I}_{s,k}^H \right) \mathbf{W}^H. \end{aligned} \quad (16)$$

Comparison of the CRB for estimating \mathbf{r}_{ave} from \mathbf{r} given by (15) with the realised covariance of \mathbf{r}_{ave} given by (16) shows that they are not the same in general. This may come as a surprise, since we are used to increasing our integration time to reduce the noise on our data. It is instructive to demonstrate that integration over the same number of STIs on all baselines is one of the possible ways to make the covariance of \mathbf{r}_{ave} equal to the CRB, thereby confirming our intuition that averaging of data to reduce the noise per data value does not affect our ability to estimate parameters from those data. If we integrate over K STIs on all baselines, we get $\mathbf{I}_{s,k} = \mathbf{I} \in \mathbb{R}^{P^2 \times P^2}$ and $\mathbf{W} = \frac{1}{K} \mathbf{I} \in \mathbb{R}^{P^2 \times P^2}$, where \mathbf{I} denotes the identity matrix of appropriate size. This simplifies the CRB to

$$\mathbf{C} = \frac{1}{K} \text{cov}(\mathbf{r}_1), \quad (17)$$

where we assumed that the covariance is the same for all STIs, and (16) to

$$\text{cov}(\mathbf{r}_{\text{ave}}) = \frac{1}{K} \text{cov}(\mathbf{r}_1), \quad (18)$$

which is equal to the CRB as expected. The factor $1/K$ confirms our common knowledge that the noise reduces with the square root of the number of samples when we integrate (remember that the CRB represents the (co)variance of the estimated parameters while noise is usually expressed in standard deviations).

Another special case, in which the covariance of \mathbf{r}_{ave} becomes (approximately) identical to the CRB, is the regime in which the instantaneous SNR is low. In this regime, the system temperature of the receiving elements (aperture array stations or dishes) is significantly higher than the source temperature. Therefore, the power measured in the auto-correlations is significantly higher than the power observed in the crosscorrelations, so we may assume that $\mathbf{R}_k \approx \sigma_n \mathbf{I}$ for an array of identical receiving elements. If all received source and noise signals can be described as independent and identically distributed Gaussian noise, $\text{cov}(\mathbf{r}_k) = \frac{1}{N} \sigma^2 \mathbf{I} \in \mathbb{C}^{P^2 \times P^2}$. This gives $\text{cov}(\mathbf{r}) = \frac{1}{N} \sigma^2 \mathbf{I} \in \mathbb{C}^{KP^2 \times KP^2}$. Substituting this in (13) gives

$$\mathbf{C} = \left(\mathbf{I}_s \left(\frac{\sigma^2}{N} \mathbf{I} \right)^{-1} \mathbf{I}_s^H \right)^{-1} = \left(\frac{N}{\sigma^2} \mathbf{I}_s \mathbf{I}_s^H \right)^{-1} = \frac{\sigma^2}{N} \mathbf{W}, \quad (19)$$

where we used (8) in the last step.

Using the same result in (16), we find

$$\text{cov}(\mathbf{r}_{\text{ave}}) = \mathbf{W} \mathbf{I}_s \frac{\sigma^2}{N} \mathbf{I}_s^H \mathbf{W} = \frac{\sigma^2}{N} \mathbf{W} \mathbf{I}_s \mathbf{I}_s^H \mathbf{W} = \frac{\sigma^2}{N} \mathbf{W}, \quad (20)$$

which is equal to the CRB derived in (19). We can therefore conclude that the BDA scheme described in Sec. 2.1 works well in cases where the instantaneous SNR, i.e., the SNR before integration, of the observed sources is very low. However, in cases in which the instantaneous SNR is close to unity or even higher, a more advanced BDA scheme is required. Unfortunately, such a BDA scheme will require significantly more computing resources. Since the sources in the high-SNR case are observed with very high SNR, it is unlikely that a small loss of information will lead to a significant reduction in the scientific output of such observations. This intuition is corroborated by the analysis presented in Sec. 2.3 that demonstrates that this information loss does not affect our imaging capability directly.

The conclusion that some information is lost when applying BDA as described in Sec. 2.1 in the high-SNR regime obviously raises the following questions:

- (i) How much information is lost? In other words, how much does the noise on the averaged visibilities increase compared to the optimal case in which no information is lost?
- (ii) Is there a BDA scheme that does not incur an information loss and, if so, what is this scheme?

These issues are addressed in the next two subsections. These subsections aim to provide some more insight into the information loss issue identified above, but do not affect the conclusions already drawn or provide details that are used later in the paper. The reader may therefore skip these two subsections if desired.

2.2.1 A statistically optimal BDA scheme

In this section, we derive a statistically efficient estimator to estimate \mathbf{r}_{ave} from \mathbf{r} . A statistically efficient estimator is unbiased and provides the lowest possible covariance on the estimated parameters. A maximum-likelihood (ML) estimator is such an estimator. It is known that optimization of the weighted least squares cost function with covariance matched weighting leads to estimates that are asymptotically, i.e., for a large number of samples, equal to ML estimates (Ottersten et al. 1998). We thus want to solve

$$\hat{\mathbf{r}}_{\text{ave}} = \underset{\mathbf{r}_{\text{ave}}}{\operatorname{argmin}} \left\| \operatorname{cov}(\mathbf{r})^{-1/2} \left(\hat{\mathbf{r}} - \mathbf{I}_s^H \mathbf{r}_{\text{ave}} \right) \right\|^2. \quad (21)$$

The solution to this problem is well known and is given by

$$\hat{\mathbf{r}}_{\text{ave}} = \left(\operatorname{cov}(\mathbf{r})^{-1/2} \mathbf{I}_s \right)^\dagger \operatorname{cov}(\mathbf{r})^{-1/2} \hat{\mathbf{r}}, \quad (22)$$

where \dagger denotes the pseudo-inverse. Using the Moore-Penrose inverse, we can write this as

$$\hat{\mathbf{r}}_{\text{ave}} = \left(\mathbf{I}_s \operatorname{cov}(\mathbf{r})^{-1} \mathbf{I}_s^H \right)^{-1} \mathbf{I}_s \operatorname{cov}(\mathbf{r})^{-1} \hat{\mathbf{r}}. \quad (23)$$

Using the standard error propagation formula, as we did before to obtain (16), we find,

$$\begin{aligned} \operatorname{cov}(\mathbf{r}_{\text{ave}}) &= \left(\mathbf{I}_s \operatorname{cov}(\mathbf{r})^{-1} \mathbf{I}_s^H \right)^{-1} \mathbf{I}_s \operatorname{cov}(\mathbf{r})^{-1} \operatorname{cov}(\mathbf{r}) \times \\ &\quad \operatorname{cov}(\mathbf{r})^{-1} \mathbf{I}_s^H \left(\mathbf{I}_s \operatorname{cov}(\mathbf{r})^{-1} \mathbf{I}_s^H \right)^{-1} \\ &= \left(\mathbf{I}_s \operatorname{cov}(\mathbf{r})^{-1} \mathbf{I}_s^H \right)^{-1}, \end{aligned} \quad (24)$$

which is equal to the CRB for estimating \mathbf{r}_{ave} from \mathbf{r} derived in (13). This shows that it is possible to develop a BDA scheme that incurs no information loss (besides the expected decorrelation discussed later). However, the required weighting is data dependent, which means that the weighting matrix cannot be precomputed, and is computationally expensive, even if the block-diagonal structure of $\operatorname{cov}(\mathbf{r})$ is exploited. Since information loss mainly occurs in high-SNR scenarios, it is probably acceptable and applying such an advanced BDA scheme is probably not worth the investment. This conclusion is examined in more detail below.

2.2.2 Quantification of information loss

To assess the impact of the instantaneous SNR on estimation performance, we conduct a simulation in which we estimated $\hat{\mathbf{r}}_{\text{ave}}$ from $\hat{\mathbf{r}}$ using the optimal estimator described by (23) as well as using the proposed BDA scheme described by (9).

For this simulation, we use a simple visibility model consisting of a single source with power $\sigma_s = 1$ in the phase center and uncorrelated additive noise with power $\sigma_n = \sigma_s / \text{SNR}$ per receiving element. For simplicity, we assume that the source model is constant over all STIs. We thus have

$$\mathbf{R}_k = \sigma_s \mathbf{1} \mathbf{1}^H + \sigma_n \mathbf{I} \in \mathbb{C}^{P \times P} \quad \forall k, \quad (25)$$

where $\mathbf{1}$ denotes a column vector filled with ones of appropriate length. Note that this model is independent of the array layout as the source is in the phase center. Assuming that the signals, from which \mathbf{R}_k is estimated, can all be described by independent and identically distributed complex

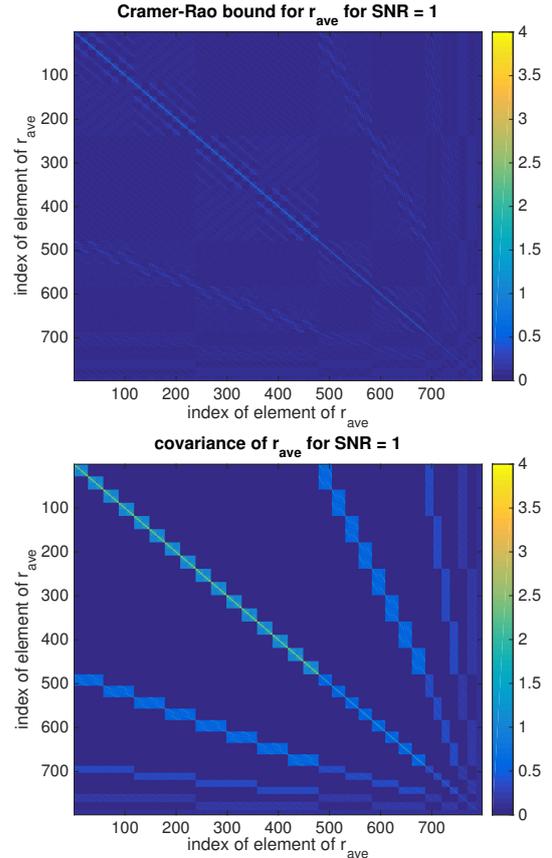


Figure 1. CRB (top) and covariance of \mathbf{r}_{ave} when estimated using the normal BDA scheme described by (9) for an SNR of unity.

Gaussian signals, we have (see, for example, appendix C in Boonstra (2005))

$$\operatorname{cov}(\mathbf{r}_k) = \frac{1}{N} \overline{\mathbf{R}}_k \otimes \mathbf{R}_k, \quad (26)$$

where overbar denotes conjugation and \otimes denotes the Kronecker product of two matrices.

For our comparison, we calculated

- (i) the CRB for estimating \mathbf{r}_{ave} from \mathbf{r} using (13);
- (ii) the covariance of \mathbf{r}_{ave} for the BDA scheme described by (9) using the standard error propagation formula as in (16);
- (iii) the covariance of \mathbf{r}_{ave} when \mathbf{r}_{ave} is estimated using the optimal BDA scheme described in (23).

In our simulations, we varied the SNR from 10^{-3} to 1 in 19 steps (20 SNR values in total) equidistantly spaced on a logarithmic scale. For clarity, we used a simple 10-element uniform linear array for this example instead of the SKA-mid configuration that will be used later in this paper. The maximum integration time was set to 16 STIs and the number of values averaged was doubled every time the baseline length was halved.

Figure 1 compares the covariance of \mathbf{r}_{ave} estimated using the non-optimal BDA scheme described by (9) with the CRB. The covariance of \mathbf{r}_{ave} using the optimal weighting described by (23) was equal to the CRB as expected and is therefore not shown separately. The covariance of \mathbf{r}_{ave} es-

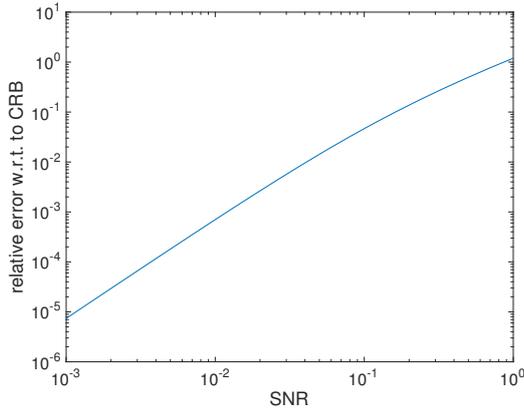


Figure 2. Mean relative increase of the variance relative to the CRB for \mathbf{r}_{ave} obtained after BDA as function of SNR as calculated using Eq. (27).

estimated using (9) is not only significantly higher than the CRB, indicating a lower estimation accuracy, but also has a different structure. This suggests a significant loss and change of information content in the data.

To assess the impact of SNR, we computed the mean relative increase of the variance relative to the CRB for \mathbf{r}_{ave} :

$$\epsilon_{\text{mean}} = \frac{1}{M} \sum_{m=1}^M \frac{|\text{[vecdiag}(\mathbf{C})]_m - \text{[vecdiag}(\text{cov}(\mathbf{r}_{\text{ave}}))]_m|}{\text{[vecdiag}(\mathbf{C})]_m}, \quad (27)$$

where $[\cdot]_m$ denotes the m th element of the vector. Figure 2 shows this measure as function of SNR. The results confirm our theoretical finding that the difference between the covariance of \mathbf{r}_{ave} obtained using BDA and the CRB only has a significant effect in high-SNR regimes. Since the sources are then observed with a high SNR, a small increase in the measurement error on the data values may be perfectly acceptable.

To get an intuition for the typical instantaneous SNR in a radio astronomical observation, let us consider an observation with LOFAR on Hercules A, one of the brightest astronomical sources in the northern hemisphere. Hercules A has a flux of 325 Jy at 178 MHz (Bennett 1962), while the source equivalent flux density (SEFD) of a single LOFAR HBA core station is about 3.5 kJy (van Haarlem et al. 2013). Based on Fig. 2, we may expect a deviation of order 1% from the CRB, which is effectively negligible given the typical contributions from other observing and data reduction errors.

2.3 Impact on imaging

In our analysis in the previous section, we found that the BDA scheme defined by (9) incurs a (small) loss of information compared to optimal estimation of \mathbf{r}_{ave} from \mathbf{r} . Unfortunately, the CRB (or FIM) of the averaged visibilities does not tell us whether the information content of the raw visibilities is preserved in optimal averaging. In this section, we address this issue by considering the ability to reconstruct an image from the raw (unaveraged) visibilities and from averaged visibilities.

The sky can be discretized in Q directions. If a plane

wave with power σ_q is impinging on the array from the q th direction, the visibility response of a perfectly calibrated array is given by (van der Veen & Wijnholds 2013)

$$\mathbf{R}_{k,q} = \sigma_q \mathbf{a}_{k,q} \mathbf{a}_{k,q}^H, \quad (28)$$

where $\mathbf{a}_{k,q} = [e^{2\pi j \boldsymbol{\xi}_{k,1} \cdot \mathbf{l}_q / \lambda}, \dots, e^{2\pi j \boldsymbol{\xi}_{k,P} \cdot \mathbf{l}_q / \lambda}]^T \in \mathbb{C}^{P \times 1}$ is the array response vector. The vector \mathbf{l}_q denotes the unit vector pointing in the q th direction and $\boldsymbol{\xi}_{k,p}$ is the position vector of the p th receiving element in the array during the k th snapshot. Vectorization of (28) and forming a superposition of all source and noise signals gives

$$\mathbf{r}_k = \left(\sum_{q=1}^Q (\bar{\mathbf{a}}_{k,q} \otimes \mathbf{a}_{k,q}) \sigma_q \right) + \text{vec}(\boldsymbol{\Sigma}_n), \quad (29)$$

where $\boldsymbol{\Sigma}_n$ denotes the noise covariance matrix. If the receiver noise powers are uncorrelated, the noise covariance matrix is diagonal, i.e., $\boldsymbol{\Sigma}_n = \text{diag}(\boldsymbol{\sigma}_n)$. Stacking the array response vectors $\mathbf{a}_{k,q}$ in an array response matrix $\mathbf{A}_k = [\mathbf{a}_{k,1}, \dots, \mathbf{a}_{k,Q}] \in \mathbb{C}^{P \times Q}$ and the source powers σ_q in a vector $\boldsymbol{\sigma}_s = [\sigma_1, \dots, \sigma_Q]^T$, we can write (29) as

$$\begin{aligned} \mathbf{r}_k &= (\bar{\mathbf{A}}_k \circ \mathbf{A}_k) \boldsymbol{\sigma}_s + (\mathbf{I} \circ \mathbf{I}) \boldsymbol{\sigma}_n \\ &= \begin{bmatrix} \bar{\mathbf{A}}_k \circ \mathbf{A}_k & \mathbf{I} \circ \mathbf{I} \end{bmatrix} \begin{bmatrix} \boldsymbol{\sigma}_s \\ \boldsymbol{\sigma}_n \end{bmatrix} \\ &= \mathbf{M}_k \boldsymbol{\theta}, \end{aligned} \quad (30)$$

where \circ denotes the Khatri-Rao or column-wise Kronecker product of two matrices. We have also introduced the parameter vector $\boldsymbol{\theta} = [\boldsymbol{\sigma}_s^T, \boldsymbol{\sigma}_n^T]^T$ and the measurement matrix $\mathbf{M}_k \in \mathbb{C}^{P^2 \times (Q+P)}$. Stacking all snapshots in a single visibility vector, we get

$$\mathbf{r} = \begin{bmatrix} \mathbf{M}_1 \\ \vdots \\ \mathbf{M}_K \end{bmatrix} \boldsymbol{\theta} = \mathbf{M} \boldsymbol{\theta}. \quad (31)$$

In our image reconstruction process, we want to solve

$$\hat{\boldsymbol{\theta}} = \underset{\boldsymbol{\theta}}{\text{argmin}} \left\| \text{cov}(\mathbf{r})^{-1/2} (\hat{\mathbf{r}} - \mathbf{M} \boldsymbol{\theta}) \right\|^2, \quad (32)$$

where we have assumed covariance matched weighting to obtain an estimate that is asymptotically equivalent to the ML estimate (Ottersten et al. 1998). The standard solution to this problem is given by

$$\begin{aligned} \hat{\boldsymbol{\theta}} &= \left(\mathbf{M}^H \text{cov}(\mathbf{r})^{-1} \mathbf{M} \right)^{-1} \mathbf{M}^H \text{cov}(\mathbf{r})^{-1} \hat{\mathbf{r}} \\ &= \mathbf{M}_{\text{im}} \hat{\mathbf{r}}, \end{aligned} \quad (33)$$

i.e., the image reconstruction process can be described by a matrix $\mathbf{M}_{\text{im}} \in \mathbb{C}^{(Q+P) \times KP^2}$ that maps the measured visibilities on the parameters describing the image and the instrumental noise. It is interesting to note that this simple equation represents the imaging process including deconvolution in the case of weighted least squares image reconstruction while it can also represent a simple DFT imager or creation of a dirty image by gridding followed by a fast Fourier transform (van der Veen & Wijnholds 2013). In other words, (33) describes a whole class of image reconstruction methods in which the visibilities can be described as a linear superposition of sources signals.

Using standard error propagation, as before, the covariance of $\boldsymbol{\theta}$ is given by

$$\begin{aligned} \text{cov}(\boldsymbol{\theta}) &= \left(\frac{\partial \boldsymbol{\theta}}{\partial \mathbf{r}^T} \right) \text{cov}(\mathbf{r}) \left(\frac{\partial \boldsymbol{\theta}}{\partial \mathbf{r}^T} \right)^H \\ &= \mathbf{M}_{\text{im}} \text{cov}(\mathbf{r}) \mathbf{M}_{\text{im}}^H. \end{aligned} \quad (34)$$

If we use BDA using a scheme commensurate with (9) and apply the image reconstruction process described by \mathbf{M}_{im} to the full resolution visibility vector restored from \mathbf{r}_{ave} using (10), we obtain the reconstructed image give by

$$\hat{\boldsymbol{\theta}}_{\text{ave}} = \mathbf{M}_{\text{im}} \mathbf{I}_s^H \hat{\mathbf{r}}_{\text{ave}}. \quad (35)$$

In this case, the covariance of the reconstructed parameters will be

$$\text{cov}(\boldsymbol{\theta}_{\text{ave}}) = \mathbf{M}_{\text{im}} \mathbf{I}_s^H \text{cov}(\mathbf{r}_{\text{ave}}) \mathbf{I}_s \mathbf{M}_{\text{im}}^H. \quad (36)$$

Substitution of (16) gives

$$\text{cov}(\boldsymbol{\theta}_{\text{ave}}) = \mathbf{M}_{\text{im}} \mathbf{I}_s^H \mathbf{W} \mathbf{I}_s \text{cov}(\mathbf{r}) \mathbf{I}_s^H \mathbf{W} \mathbf{I}_s \mathbf{M}_{\text{im}}^H. \quad (37)$$

In this derivation, we reconstruct the non-averaged visibilities using (10). This operation simply assigns each value in \mathbf{r}_{ave} to the entries of the reconstructed non-averaged visibility vector $\tilde{\mathbf{r}}$ corresponding to the entries of \mathbf{r} that were averaged to obtain that value in \mathbf{r}_{ave} . If \mathbf{r} describes a visibility *model* commensurate with BDA, i.e., if \mathbf{r} contains no noise and entries that are averaged in BDA have the same value, we will have $\tilde{\mathbf{r}} = \mathbf{r}$. Since $\mathbf{r}_{\text{ave}} = \mathbf{W} \mathbf{I}_s \mathbf{r}$, this implies that

$$\tilde{\mathbf{r}} = \mathbf{r} = \mathbf{I}_s^H \mathbf{W} \mathbf{I}_s \mathbf{r}. \quad (38)$$

This shows that multiplication with $\mathbf{I}_s^H \mathbf{W} \mathbf{I}_s$ has the same effect as multiplication with the identity matrix for noise-free vectors, such as vectors containing expected values or model values, that have a structure commensurate with BDA. If each column of \mathbf{M}_{im}^H has that property, it thus follows that

$$\mathbf{I}_s^H \mathbf{W} \mathbf{I}_s \mathbf{M}_{\text{im}}^H = \mathbf{M}_{\text{im}}^H. \quad (39)$$

In this case, it is easy to show that

$$\text{cov}(\boldsymbol{\theta}_{\text{ave}}) = \text{cov}(\boldsymbol{\theta}), \quad (40)$$

i.e., that the covariance of the parameters reconstructed from \mathbf{r} and \mathbf{r}_{ave} are the same, which implies that reconstruction is equally (un)successful with and without application of BDA.

To re-iterate, BDA does not affect our ability to reconstruct an image from the visibility data as long as

- (i) the relation between the source and noise signals and the visibilities can be described by a linear transformation;
- (ii) the image reconstruction process can be described by a linear transformation \mathbf{M}_{im} whose rows have a structure commensurate with BDA.

It can be easily demonstrated that the latter condition holds for ML image reconstruction using weighted least squares optimisation as discussed above and for formation of dirty images by means of a discrete Fourier transform. The derivation presented in this section holds for a perfectly calibrated array. Calibration usually requires a non-linear model. As a result, the reasoning in this section does not hold. Fortunately, most radio astronomical observations are low-SNR

observations making the information loss due to BDA negligible. Also, many radio astronomical systems are very stable requiring only small corrections, which minimises the impact of the non-linearities.

2.4 Decorrelation

Decorrelation, or amplitude loss, due to time averaging over total integration time T on the baseline with mid-point (u_0, v_0, w_0) for a source located at direction cosine (l, m, n) can be described by the amplitude reduction factor given by Eq. (18-31) in (Bridle & Schwab 1999)

$$\begin{aligned} R_T &= \\ &= \text{sinc} \left\{ \pi T \left(\left. \frac{\partial u}{\partial t} \right|_{u_0} l + \left. \frac{\partial v}{\partial t} \right|_{v_0} m + \left. \frac{\partial w}{\partial t} \right|_{w_0} (n-1) \right) \right\} \\ &\approx 1 - \frac{\pi^2 T^2}{6} \left(\left. \frac{\partial u}{\partial t} \right|_{u_0} l + \left. \frac{\partial v}{\partial t} \right|_{v_0} m \right)^2, \end{aligned} \quad (41)$$

where the approximation assumes that $\nu T \ll 1$, and that the direction cosine offset angles are small, i.e., that the decorrelation is small.

If L_x, L_y and L_z denote the baseline components along the principal axes of the ITRF coordinate system, Bridle & Schwab (1999) show that

$$\begin{aligned} \frac{\partial u}{\partial t} &= \frac{1}{\lambda} (L_x \cos H - L_y \sin H) \omega_E \\ \frac{\partial v}{\partial t} &= \frac{1}{\lambda} (L_x \sin \delta \sin H + L_y \sin \delta \cos H) \omega_E, \end{aligned} \quad (42)$$

where δ and H denote the declination and hour angle of the source and $\omega_E = 7.2925 \cdot 10^{-5}$ rad/s is the angular velocity of the Earth.

In the simulations presented in Sec. 4, we will assess the impact of BDA on imaging observations. Natural weighting provides the best sensitivity, in particular for centrally condensed array configurations, while uniform weighting provides a better angular resolution. We will therefore present results for both natural weighting and uniform weighting. For natural weighting, the expected amplitude reduction in the image can be determined by taking the average amplitude reduction factor for all baselines, where the amplitude reduction factor for each baseline is calculated by Eq. (41). For uniform weighting, we need to take into account an appropriate weighting factor per baseline. As this weighting factor depends on the density of (u, v, w) -samples, this requires calculation of $\partial w / \partial t$ as well. It is straightforward to derive this from Eq. (18-32) in (Bridle & Schwab 1999), resulting in

$$\frac{\partial w}{\partial t} = -\frac{1}{\lambda} (L_x \cos \delta \sin H + L_y \cos \delta \cos H) \omega_E. \quad (43)$$

The weighting factor for uniform weighting can now be described by

$$W_{\text{uni}} = \frac{1}{N} \sqrt{\left(\frac{\partial u}{\partial t} \right)^2 + \left(\frac{\partial v}{\partial t} \right)^2 + w_{\text{factor}} \left(\frac{\partial w}{\partial t} \right)^2} \quad (44)$$

where N is the number of baseline data values within a given (u, v, w) distance around the location of the baseline under consideration. In his Appendix E, Briggs (1995) shows that the maximum (u, v, w) distance can be specified by $\pi^{-0.5} \Delta$

where Δ would be the corresponding gridded (u, v) -plane increment. For a source 2° from the field centre this maximum distance works out to be about 8 wavelengths. Test simulation runs showed that only 4% of the simulated visibilities had N greater than 1, and just setting $N = 1$ gave minimal differences in the attenuation factors derived in the next section, but sped up the computations enormously.

At small distances from the field centre the $\partial w/\partial t$ term in the above equation basically contributes nothing but at large distances from the field centre begins to provide a contribution due to the rotation matrices necessary to shift the phase centre to the new location. Empirical tests (see Sec. 4.4) showed that giving the w_{factor} term a value of about 0.15 gives a good fit to the data.

BDA can also be applied over frequency. Bridle & Schwab (1999) show that the effect of frequency smearing can be described as a convolution in the image domain with a distortion function. This distortion function causes a radial broadening of the point source response. Although the integrated flux density of the source is preserved, this broadening causes a reduction of the amplitude proportional to $1/\left(\Delta\nu\sqrt{l_0^2 + m_0^2}\right)$. While it would be interesting to include frequency averaging in our simulations, simulated SKA observations at just a single frequency already required significant computing resources. So we will only investigate the effects of time averaging in the simulations described in the next section and therefore not discuss the effects of BDA over frequency in more detail here.

3 BDA AND CALIBRATION

A practical issue with visibility data obtained using BDA is that visibility samples may have different integration times and, as a result, that the time series for different baselines have different lengths. Calibration routines will thus have to carry out appropriate interpolations to account for the different sampling times and appropriate weighting to account for the different integration times. Not all calibration routines may be able to do this. Another possible calibration issue is that the time scales of certain instrumental and environmental variations are such that these variations are well sampled by the longest baselines (with the shortest integration times) while these time scales may be comparable with the longest integration time on the shortest baselines. Both issues can be handled by the Compress-Expand-Compress (CEC) method proposed by Salvini & Wijnholds (2017). As calibration is a major concern for modern radio interferometers, we summarize this method and recapitulate the main results from Salvini & Wijnholds (2017) below.

The CEC method consists of three steps:

Compression: After crosscorrelation and, if desired, flagging, BDA is applied following (9). This ensures that the amount of data that needs to be transferred to the compute cluster performing calibration (and likely imaging) is reduced.

Expansion: The averaged visibilities are then used to fill the time series of the corresponding visibilities at the full data rate as described by (10). Calibration can then be carried out at the full time resolution.

Compression: After calibration, the data are compressed

again using the desired BDA scheme before imaging. This compression step reduces the amount of data that needs to be gridded. As the short-term variations have been corrected by calibration, the compression at this stage could be even higher than during the first compression stage.

As calibration and imaging may be done on the same computing platform, the benefits of the second compression stage may not be obvious. However, as the first stage of calibration is necessarily either direction-independent calibration or calibration in a very limited number of directions due to the low SNR of the sources observed, we can safely assume that the compute requirements for this initial calibration step are far less than the compute requirements for gridding. The second compression step may thus lead to a significant reduction of computing resources needed for gridding.

A possible concern for the first compression stage is that the temporal variation of the visibility values on, in particular, the short baselines within the averaging interval may not only be caused by variations in the geometrical delay, but also by temporal variations in some calibration parameters. Salvini & Wijnholds (2017) showed that this risk can be mitigated by using not only mid-point averaging (the zeroth order moment) for BDA, but also the first, and possibly even higher, order moments of the visibility data over the averaging interval at the expense of lower data volume reduction during the first compression stage. Salvini & Wijnholds (2017) demonstrated with a simulation with extreme instrumental gain variations that the use of first (and second) order moments allows significant compression while retaining the ability to reconstruct $1\text{-}\mu\text{Jy}$ sources in the vicinity of a 1-Jy source, indicating that over 60 dB dynamic range remains feasible while using BDA in the presence of fast gain variations.

Based on these findings, we decided to focus the simulations presented in Sec. 4 on validating the conclusion drawn based on the theoretical analysis presented in Sec. 2 that BDA has no detrimental effect on imaging performance besides a predictable level of decorrelation. This allowed us to set up the simulations in such a way that this particular aspect could be isolated from other effects that need to be dealt with when implementing a data processing pipeline for an actual system.

4 SIMULATION OF TIME AVERAGING EFFECTS

4.1 Description of simulations

In addition to the theoretical work presented earlier we made simulated observations with the SKA1-mid telescope configuration to check if BDA had direct effect on decorrelation of "actual" observations. We used the SKA1-mid antenna positions as provided by Heystek (2015) plus the locations of the 64 MeerKAT (Jonas 2009) antennas for a total of 197 antennas. With the given antenna positions, the actual physical baseline lengths ranged from 157 km down to 29 m.

We made single channel observations covering the declination range of 10° to -70° in 20° steps for a total of five different declinations. At each declination we observed seven short cut observations of 7 minutes each with an initial cor-

relator dump time of 0.14 seconds on all of the 19306 baselines. The short cut observations were made at hour angles -4.5, -3, -1.5, 0, 1, 2 and 4 hours. Each short observation was simulated and processed separately. The observing frequency was 700 MHz.

The simulation process was as follows:

- Place a 1-Jy point source at an offset from the field centre of 2° along either direction cosine m (roughly declination) or l (roughly right ascension). Note that an offset of 2° at 700 MHz corresponds to an equivalent offset of 1° at 1400 MHz. We expect these distances to correspond roughly to the 10 percent power response of the dish primary beam and thus to be about the maximum distance from the field centre at which one might expect to do "significant" science.

- Generate noise-free visibilities with the above mentioned antenna locations and source position using the CASA simulation tool.

- As the baselines become shorter, average the sampled visibilities according to various criteria (see below) and replace the original "raw" visibilities by their "average" at the same (u, v) location. This allows us to minimize any errors due to different (u, v) coverage. Also, the CASA-based imager that we used, assumes that the WEIGHTS column of a Measurement Set is an indicator of system temperature changes and is not an indicator for how much weight should be assigned to a (u, v) point at the gridding stage of imaging. Obviously, in an actual SKA system, the (u, v) data would need to be appropriately weighted to represent the (u, v) coverage each averaged point represents.

- Make images from the averaged visibilities using the extremes of natural and uniform weighting (Briggs 1995). Natural weighting essentially weights the (u, v) grid proportional to the signal collected in a given area whereas uniform weighting essentially normalizes the signal in each area. In the current SKA1 design, there are many more visibilities at short spacings than long ones, so natural weighting emphasizes short spacings.

- Determine the difference between the maximum value found in the images and the expected 1-Jy signal as both a function of hour angle and declination. We could keep the size of the images small (1024×1024 pixels) by phase-rotating the reference position of the image to the position of the offset target source.

Note that in the following plots signals that are minimized at hour angle zero are due to the test source offset in direction cosine l while signals that are maximized at hour angle zero are due to the test source offset in direction cosine m . The 6-hour phase offset in time is not surprising as the sources are offset in the sky by exactly 90 degrees.

4.2 Effects due to correlator integration

The current SKA1 design calls for the correlator to collect data for 0.14 seconds before dumping the result. We find that even this short integration time could have a significant effect on imaging. We determined the magnitude of this effect by taking the (u, v, w) locations as calculated at 0.14-s intervals, linearly interpolating them into 20 subintervals, calculating the expected amplitudes and phases at the subintervals and then averaging them together. The results are shown in Fig. 3. In addition to the results of the

simulated observations, the plots also show (solid lines) a prediction of the expected decorrelation for uniform weighting as a function of hour angle and declination based on the prescription provided in Sec. 2.4.

We checked that these results made sense by actually making some simulations where we lowered the correlator dump time to 1/20th of 0.14 s or 0.007 s and then averaged up the signals in 0.14-s blocks on all baselines. We got very similar results. While a residual of 0.25% sounds small, 0.25% of a 10-mJy source equates to $25 \mu\text{Jy}$, a signal expected to be well above the SKA signal sensitivity limit.

4.3 A simple averaging scheme

Since our physical baseline length ranges from 157 km down to 29 m, a simple way to average is to double the integration time for every factor 2 decrease in baseline length. This leads to the following baselines zone limits and associated baseline averaging:

physical baseline zones (km)
 { 80, 40, 20, 10, 5, 2.5, 1.25, 0.625, 0.3125 }
 baseline averaging (multiples of correlator dump time)
 { 1, 2, 4, 8, 16, 32, 64, 128, 256, 512 },

which means that baselines longer than 80 km are not averaged, baselines with length between 40 and 80 km have two points averaged, baselines with length between 20 and 40 km have four points averaged, and so on.

This scheme gave an 87.5% reduction in data. However we found that the following scheme:

physical baseline zones (km)
 { 80, 40, 30, 20, 15, 10, 7.5, 5, 3.75, 2.5, 1.875, 1.25, 0.9375, 0.625, 0.5625, 0.375, 0.28125 }
 baseline averaging (multiples of correlator dump time)
 { 1, 2, 3, 4, 6, 8, 12, 16, 24, 32, 48, 64, 96, 128, 192, 256, 384, 512 },

where we inserted some intermediate binning intervals, gave almost as large a reduction in visibility data volume, 86.8%, but decreased the residual errors. For example, the maximum residual error decreased from 0.73% to 0.64% at hour angle zero for the source offset 2° in direction cosine m at declination -70° (see Fig. 4). We will refer to this BDA scheme as Scheme 1.

In this analysis, we used actual physical baseline lengths rather than projected (u, v) lengths. The maximum residuals as a function of hour angle and declination for a test source offset by 2 degrees in direction cosines l and m are shown in Fig. 4. Note that the observed values and theoretical predictions (solid curves) shown in the figures implicitly include the effects of the 0.14-s correlator integration time.

Despite the fact that we have been averaging data from a large number of baselines with a wide variety of orientations, the behaviour of the residuals shown in Fig. 3 and Fig. 4 is astonishingly consistent with the simple analysis of time-averaging given on pages 206 through 208 of (Thompson et al. 1998). At extreme hour angles, the corrugations of interferometer fringes are parallel to the l -direction cosine but the Earth's rotation vector direction is perpendicular to the l -direction and l -fringe corrugations. So a source located along the l -direction cosine will be moving through fringes most rapidly and be most susceptible to averaging effects,

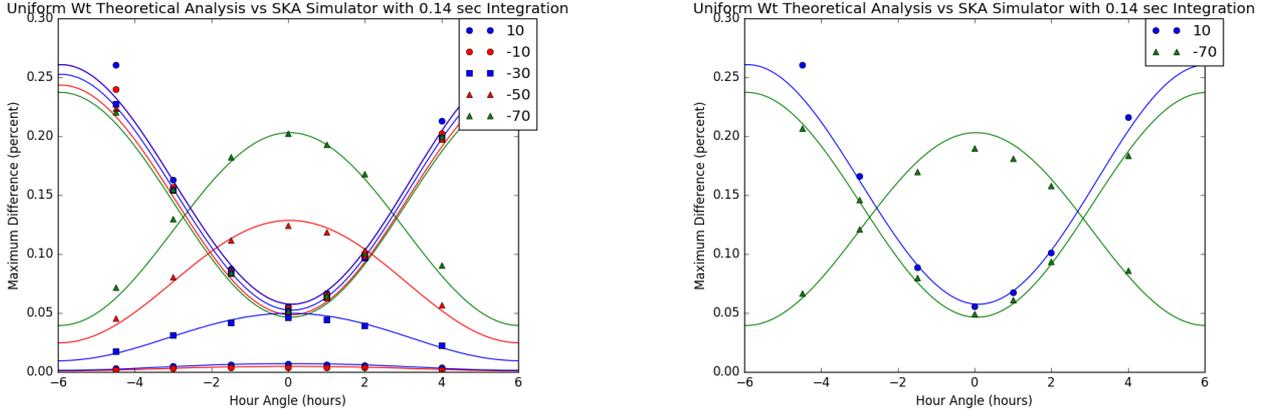


Figure 3. These plots show decorrelation effects on 1-Jy point sources situated 2° from the phase reference position at 700 MHz due to a correlator dump time of 0.14 s. In these and succeeding plots, the residual errors peak near hour angle 0 for a source placed at an offset of 2° in direction cosine m and the errors are minimized near hour angle 0 for a source located at an offset of 2° in direction cosine l . The actually measured data points are shown as separate symbols, while the solid lines represent the theoretically predicted attenuation based on the prescription in Sec. 2.4. We obtained these results by dividing each original (u, v) step into 20 subintervals and then averaging the calculated amplitudes and phases over the 20 subintervals. While the effects are very small for natural weighting (less than about 0.04 percent) they increase to a maximum of about 0.25 percent for uniform weighting. The left plot shows the maximum residuals in images made with uniform weighting as a function of hour angle and declination. The right plot shows the equivalent results found when we generated actual (u, v) positions with a correlator sample time of 0.007 s and then averaged the data together in blocks of 0.14 s. To obtain the right plot we only calculated the results expected for the most extreme variations as the simulations took about 15 hours of computer time for each data point. For this reason, we did not bother to do simulations at 0.007s resolution for the source offset in direction cosine m at 10° declination. The errors are already very small as shown in the left plot. Decorrelation is expressed as the fractional reduction of power (measured as difference between the original power and observed power) in the original 1-Jy signal. The two approaches give good agreement and validate our approach of linearly interpolating between original (u, v) points to simulate the correlator effects.

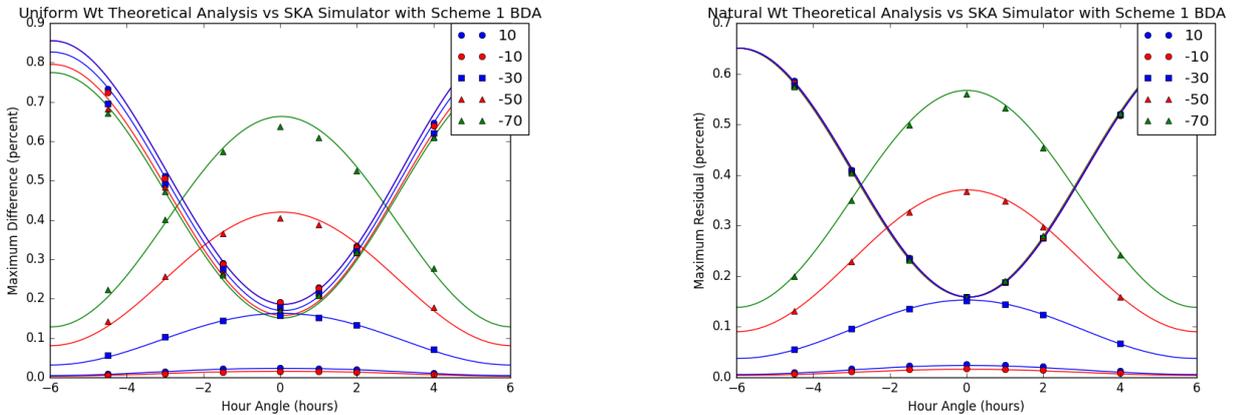


Figure 4. These plots show results for a 1-Jy point source placed at offsets of 2° in direction cosine m (residual errors peak at hour angle 0) and at 2° in direction cosine l (residual errors are minimized at hour angle 0). In the left plot we show the maximum residuals in images made with uniform weighting as a function of hour angle and declination. The right plot shows the equivalent results from images made with natural weighting. Note that the behaviour is very similar. Decorrelation is expressed as a percentage of the original test source signal of 1 Jy.

whereas the Earth's rotation causes a source along the m -direction to be moving approximately parallel to fringe corrugations so averaging has little effect. The inverse situation occurs at transit. Thompson et al. (1998) also show (their equation 6.81) that residual effects in m should be subject to a $\sin^2(\text{dec})$ effect, and indeed the increase in m residuals at transit follows quite closely $\sin^2(\text{dec})$ behaviour.

Although Thompson et al. (1998) did not derive any

declination dependent effects for a source offset in direction cosine l our results do show a small effect. We believe that this is due to the fact that, in order to make small 1024×1024 pixel images, we have to do a rotation of the phase centre to the position of the offset source. This requires a sequence of rotation matrices which will cause the w -coordinate to start contributing to the u and v velocity terms.

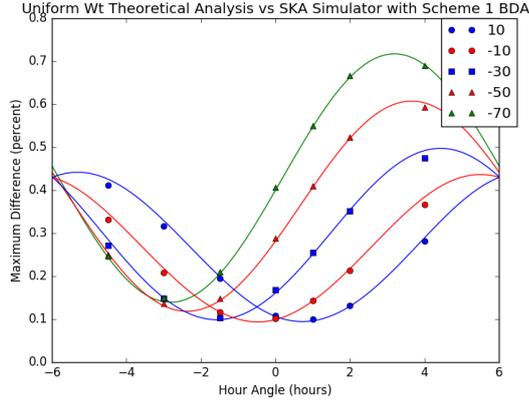


Figure 5. This figure shows the predicted and observed decorrelation as a function of hour angle when the test source is placed at an offset of 1.414 degrees in both l and m direction coordinates.

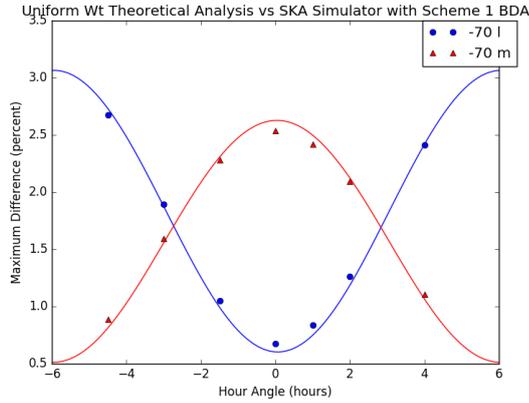


Figure 6. This figure shows the predicted and observed decorrelation as a function of hour angle for Scheme 1 BDA but with double integration time on all baselines. We get rid of 93% of the data (from 87% in the original scheme) while the maximum residual error increases to about 2.6% from about 0.65% previously.

The plots shown so far show a symmetry with respect to hour angle. However be reminded that this is simply due to the special positions at which we have placed the test source. If we put the test source at an offset of 1.414 degrees in both l and m direction coordinates, we get an asymmetric effect as shown in Fig. 5. Essentially every source in the field will follow a different pattern.

By using Scheme 1 BDA we managed to shrink the amount of data by 87%. If we double the averaging period on all baselines (so even on the longest baselines we now average two points together) we can even reduce the amount of visibility data by 93% at the expense of increasing the residual errors to 2.6% (see Fig. 6). This error may still be acceptable for SKA data processing at large distances from the field centre.

Another interesting effect is that even the nominal correlator dump time must be included in predictions related to BDA. In Fig. 7 we show the effect of having inverted a 1-Jy source into the (u, v) -plane and applying our BDA scheme but without adjusting for the internal 0.14 s integration in-

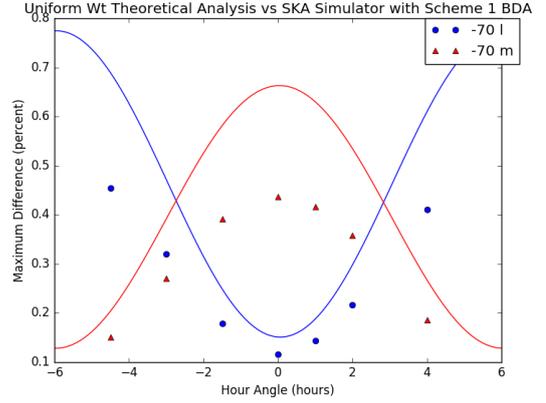


Figure 7. This figure shows the predicted and observed decorrelation as a function of hour angle when a source is inverted into the (u, v) -plane and BDA is done but the 0.14-s nominal correlator dump time is ignored.

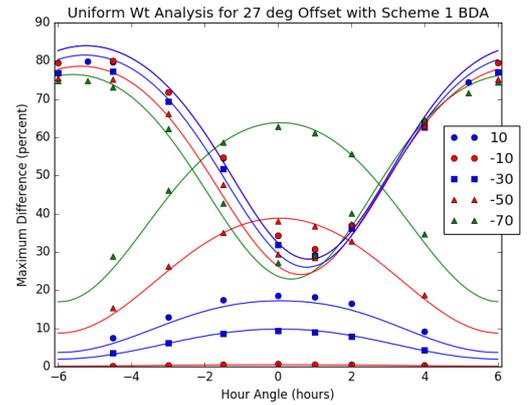


Figure 8. This figure shows the predicted and observed decorrelation as a function of hour angle and declination when uniform weighting is used to image a source situated 27.2° from the field centre.

side the correlator. It is clear that the effects of BDA are underestimated.

4.4 Wide Field Effects

Our previous simulations were restricted to sources within a few degrees of the field centre. However, as radio interferometer sensitivity increases, there is an increasing possibility that strong background sources will appear above the background noise out to great distances from the central field of view. Out of curiosity we placed a source at 27.2° from the field centre (so roughly 13 primary beams at 700 MHz) and applied the BDA averaging scheme described in this section to the visibilities. At this distance we have to use the full sinc decorrelation function given in the first part of Eq. (41) when calculating the expected decorrelation. The results are shown in Fig. 8. The actual simulated observations show excellent agreement with the theoretical predictions and suggest that the techniques discussed here could be adapted to remove the effects of such sources from actual observations.

4.5 Second averaging scheme

A potential disadvantage of BDA Scheme 1 proposed above is that doubling the integration time every time the baseline length halves may cause sharp changes in the level of decorrelation at specific spatial scales. To smooth the variation of decorrelation with spatial scale even further, we developed a second scheme. This BDA Scheme 2 works as follows: define a maximum allowable number of points to average. Then the averaging value for a given baseline becomes the minimum of the maximum allowable value and the integer value of the maximum baseline length divided by the current baseline length. So we average baselines up to a maximum allowable value and all shorter baselines are averaged to this maximum value. This scheme was used by Salvini & Wijnholds (2017) to test the CEC method proposed in that paper and summarized in Sec. 3. We varied the maximum allowable value from 32 to 500. The corresponding amount of data compression ranged from 87.8% to 89.1%, so there was rather little gain in allowing for very long integration times on the shortest baselines. In Fig. 9 we show the results for natural weighting obtained with this scheme, which confirm the earlier conclusions drawn based on our simulations with BDA Scheme 1.

One thing that is very noticeable in these results is that the residuals increase as the maximum compression value increases. The natural weighting residuals for a maximum averaging value of 32 agree well with those found for the first BDA scheme but become worse as the maximum averaging value increases as can be seen by comparing the plots from Figure 9 with the right hand plot of Figure 4.

The BDA schemes described above reduce the visibility data volume by at least 87%. We tested different averaging schemes and found that schemes with a similar level of decorrelation give a similar reduction in data volume. We suggest that BDA can significantly decrease the amount of SKA data that would need to be sent through SKA gridding and imaging pipelines. The residual errors of the suggested schemes are less than one percent at the edge of the primary beam field-of-view. An advantage of using an averaging scheme based on the physical baseline length is that for extensive surveys covering a large region of sky the data are processed in a uniformly consistent way.

5 CONCLUSIONS

In this paper, we presented a theoretical analysis indicating that baseline dependent averaging (BDA) does not impair the imaging capabilities of the instrument. If significant calibration corrections need to be made, some information is lost, but our analysis indicates that this only affects observations with a high instantaneous SNR, i.e., observations of sources whose source temperature is comparable with the system temperature. In such cases, this loss is probably acceptable.

These theoretical findings were confirmed by simulations, which did not show any indication of other detrimental effects than well-understood decorrelation. Our simulations indicate that by using BDA over time only, the visibility data volume of the SKA1-mid telescope can be reduced by over 87% while keeping the decorrelation of sources at the

edge of the field-of-view below 1%. Even more aggressive visibility data reduction could be achieved for continuum observations by applying BDA along the frequency dimension as well.

6 ACKNOWLEDGEMENT

The authors would like to thank Jeroen Stil, Richard Dodson, Daniel Mitchell, Benjamin Mort and Fred Dulwich for the useful suggestions and comments they made while we were working on the material presented in this paper.

REFERENCES

- Atemkeng M. T., Smirnov O. M., Tasse C., Foster G., Jonas J., 2016, *Monthly Notices of the Royal Astronomical Society*, 462, 2542
- Bennett A., 1962, *Memoirs of the Royal Astronomical Society*, 68, 163
- Boonstra A. J., 2005, PhD thesis, Delft University of Technology, Delft (The Netherlands)
- Bridle A. H., Schwab F. R., 1999, in Taylor G. B., Carilli C. L., Perley R. A., eds, *ASP Conference Series*, Vol. 180, *Synthesis Imaging in Radio Astronomy II*. *Astronomical Society of the Pacific*, Chapt. 18, pp 371–382
- Briggs D. S., 1995, PhD thesis, The New Mexico Institute of Mining and Technology, Socorro, New Mexico
- Cotton W. D., 2009, Technical report, *Effects of Baseline Dependent Time Averaging of UV Data*. National Radio Astronomy Observatory
- Dewdney P. E., Hall P. J., Schilizzi R. T., Lazio T. J. L. W., 2009, *Proceedings of the IEEE*, 97, 1482
- Heystek L., 2015, Technical Report SKA-TEL-INSA-0000537, SKA1_MID physical configuration coordinates. SKA Office
- Jonas J. L., 2009, *Proceedings of the IEEE*, 97, 1522
- Kay S. M., 1993, *Fundamentals of Statistical Signal Processing - Volume I: Estimation Theory*. Prentice Hall Signal Processing Series, Prentice Hall
- Mitchell D. A., et al., 2008, *IEEE Journal of Selected Topics in Signal Processing*, 2, 707
- Ottersten B., Stoica P., Roy R., 1998, *Digital Signal Processing, A Review Journal*, 8, 195
- Perley R., et al., 2009, *Proceedings of the IEEE*, 97, 1448
- Salvini S., Wijnholds S. J., 2017, in XXXIIth URSI General Assembly and Scientific Symposium (GASS) High-compression Baseline Dependent Averaging. Montreal, Canada
- Skipper C., 2014, Technical report, *Time and channel averaging*. University of Southampton
- Thompson A. R., Moran J. M., Swenson G. W., 1998, *Interferometry and Synthesis in Radio Astronomy*. Krieger Publishing Company, Malabar (Florida)
- Tingay S. J., et al., 2013, *Publications of the Astronomical Society of Australia*, 30, 1
- van der Veen A. J., Wijnholds S. J., 2013, in Bhattacharyya S. S., Deprettere E. F., Leupers R., J. T., eds, *Handbook of Signal Processing Systems*, second edn, Springer, pp 421–464

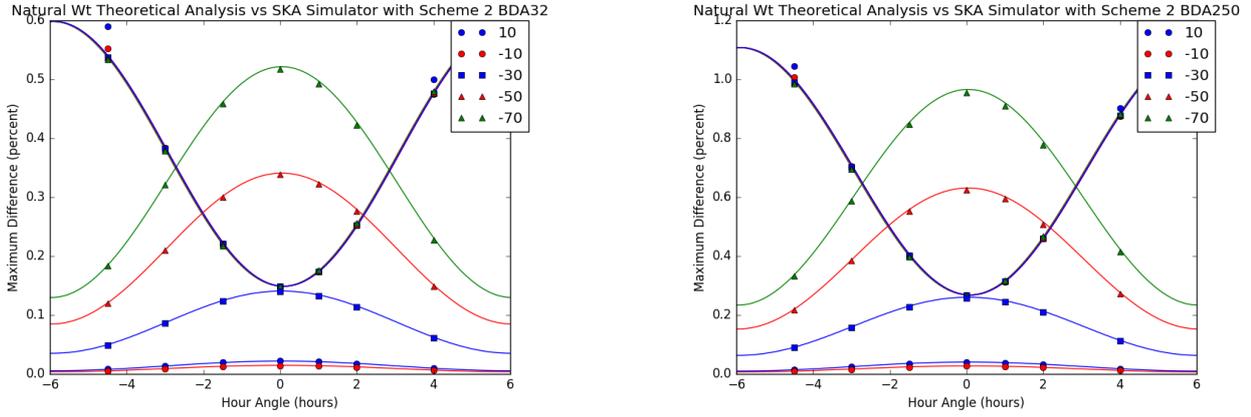


Figure 9. These plots show results for BDA Scheme 2. The plots show the variation as a function of hour angle and declination for a unit point source placed at an offset of 2° in direction cosine l and m from the field centre. The left plot shows the maximum residual in the images made with natural weighting and a maximum averaging value of 32. The right plot shows the corresponding residual behaviour for a maximum averaging value of 250.

van Haarlem M. P., et al., 2013, *Astronomy & Astrophysics*, 556, 1