

Next Generation Very Large Array Memo No. 75

Image Corruption From Antenna Pointing Errors: Simulation Results (and Useful Teaching Tool)

Christopher A. Hales (NRAO Socorro)

April 24, 2020

Abstract

I perform simulations of the ngVLA's Short Baseline Array and Main Array to assess the impacts of antenna pointing errors on image dynamic range and image fidelity. The simulations involve constructing mis-pointed antenna voltage patterns, from which corrupted baseline power patterns are obtained and used to predict corrupted visibilities. The simulations do not incorporate thermal noise or any errors other than those arising from antenna pointing. I find consistency between theoretically predicted image dynamic range limits and those measured from the simulations. I find that the ratio between antenna 2D radial pointing error (mean of Rayleigh distribution) and antenna half-power beam-width must be less than $1/13$ to yield image fidelity $F > 0.95$. The scripts developed for this memo have been made publicly available to support development of more advanced simulations that incorporate additional error terms, and to potentially serve as a teaching aid for radio astronomy.

1 Introduction

A critical design parameter for ngVLA antennas is the required pointing tracking precision. If set too small, the antenna structures will require significant stiffening to overcome environmental forces (e.g. wind loading) and the cost of the array will increase dramatically. If set too large, resultant antenna-based amplitude and phase errors will excessively corrupt the measured visibilities and in turn limit image dynamic range, inhibiting scientific utility of the array.

The purpose of this memo is twofold. The first is to verify using simulations the theoretically predicted relationship between image dynamic range and calibration errors from Perley (1999) (see also Hales 2019), where for this memo the errors are constrained to arise solely from antenna pointing errors. The second is to provide working code that can simulate interferometric observations in the presence of antenna pointing errors, upon which more sophisticated ngVLA simulations can be built. For example, future simulations could incorporate additional error terms like atmospheric delay (e.g. Pety et al. 2001) and examine the influence of these

errors on imaging performance with different subarray combinations (e.g. Tsutsumi et al. 2004; Mason 2019).

The scripts used to obtain the results presented in this memo are available at:
<https://github.com/chrishales/ngVLA-pnt-err-img-sims>

These scripts may also be useful as a teaching tool for radio astronomy. The scripts demonstrate a full simulation path from input sky model through antenna aperture illumination functions to visibilities, with flexibility to modify antenna array configuration, correlator setup, and to incorporate as many new effects as desired.

2 Image Dynamic Range Limit Arising From Pointing Errors

The image dynamic range limit arising from antenna pointing errors can be predicted theoretically as follows. First, consider the idealized relationship between the dynamic range limit D for a Stokes image (e.g. Stokes I) and antenna-based amplitude errors ϵ (correlated for both polarizations on a given antenna) for an array comprising N antennas and an observation that effectively comprises M statistically-independent snapshots (where M depends on the dynamical timescale of antenna pointing variations as well as the timescale for rotation of each baseline), given by (Perley 1999; Hales 2019)

$$D \approx \frac{N\sqrt{M}}{\epsilon}. \quad (1)$$

Next, by approximating the true profile of an antenna's primary beam (e.g. Airy disk or cosine-squared) with a Gaussian, the relationship between mean 2D radial pointing error σ_2 and antenna amplitude error ϵ resulting from observation of a source offset by angle θ from the field center can be estimated as

$$\epsilon \approx \exp[-4\alpha^2 \ln 2] - \exp[-4(\alpha + \kappa\beta)^2 \ln 2], \quad (2)$$

where $\beta = \sigma_2/\text{HPBW}$ is the fractional 2D error in units of primary beam half-power beam-width (HPBW), $\alpha = \theta/\text{HPBW}$, and worst-case projection of a 2D pointing error onto θ over multiple sources of emission yields $\kappa \approx 1$. To be clear, σ_2 is the mean of a Rayleigh distribution arising from tracking with two assumed independent axes, each characterized in 1D by zero-mean and normally distributed tracking error σ . For example, for drive motions in elevation X and azimuth Y with respective 1D errors given by σ_X and σ_Y , and assuming that any systematic pointing offsets in X and Y have been removed through pointing calibration such that $\langle X \rangle = \langle Y \rangle = 0$, the residual 2D tracking offset is described by $\rho = \sqrt{X^2 + Y^2}$ with $\sigma_X = \sigma_Y = \sigma$ and the mean 2D radial pointing error is $\sigma_2 \equiv \langle \rho \rangle = \sigma\sqrt{\pi/2}$.

Equation 2 is appropriate for a single pointing image. For a hexagonal mosaic with pointing separation $\text{HPBW}/\sqrt{2}$ (typical widest separation sufficient to obtain approximately uniform mosaic sensitivity), and assuming that each contributing primary beam is only imaged out to the HPBW (imaging further provides negligible sensitivity gain), the worst-case amplitude error in the mosaic will be located equidistant from 3 contributing pointings at their overlapping -2 dB primary beam contours (i.e. $\alpha = 1/\sqrt{6}$). The worst-case dynamic range limit in a hexagonal mosaic can therefore be estimated using Equation 1 with the amplitude error in Equation 2 reduced by factor $\sqrt{3}$ and setting $\alpha = 1/\sqrt{6}$.

3 Simulation Code Overview

Simulation code was developed in Python 3.6 with CASA 6.0¹ to analyze the effects of pointing errors on images resulting from observations with an interferometric array. Similar studies are presented by Bhatnagar et al. (2004) for EVLA and Kundert et al. (2017) for ALMA. Code from the latter (available at CASR-471) was used as a starting point for the more sophisticated simulations presented here.

The general procedure involves constructing mis-pointed antenna voltage patterns, then corrupted baseline primary beams. For each baseline, the corrupted baseline power pattern is multiplied with the true sky model and used to predict a corrupted visibility. Corrupted visibilities are constructed in this manner per polarization, baseline, frequency channel, and integration. The simulations account for both amplitude and phase errors arising from pointing errors. Imaging and deconvolution then follow. Two metrics are used to characterize image corruption: image dynamic range and image fidelity. These metrics are detailed in Section 3.1.

Following this approach, two simulation scripts were produced to analyze observations with Versions C of the ngVLA’s Short Baseline Array (SBA) and Main Array (MA). These scripts are largely similar. Differences are highlighted below, mostly regarding imaging. Specific simulation initialization parameters and resulting image corruption metrics for the SBA and MA are presented in Section 4.

General details and assumptions are as follows:

- Only pointing errors are injected. No noise or other errors are included, i.e. no bandwidth or time average smearing, no tropospheric amplitude or phase fluctuations, etc.
- For a nominated mean 2D radial pointing error σ_2 , the simulation appropriately samples normally-distributed random pointing errors in each of two dimensions per antenna, per coherence timescale. The pointing errors are assumed to be residual offsets after any slowly-varying systematic terms have been removed using pointing calibration. The scripts assume statistically independent pointing errors per integration, i.e. the pointing error coherence timescale is less than or equal to the integration time. Pointing errors are injected in the $X - Y$ linear-polarization antenna frame and are subject to parallactic angle rotation over a long observation.
- Observations centered about zero hour angle to avoid shadowing with transit 10° south of zenith
- Stokes I analysis only
- Flexibility to simulate measurement sets with multiple spectral channels (with assumed placement within a single spectral window) and multiple mosaic pointings. Individual channels have idealized zero bandwidth due to the neglect of bandwidth and time average smearing. The SBA code currently supports 3 pointing hexagonal mosaics with the field center located equidistant from the pointing centers, and 7 pointing hexagonal mosaics with the field centered at the central pointing. The MA code only supports a single pointing.
- All antennas are assumed to have identical apertures that are circular, unblocked, and with zero squint
- Rather than Fourier transform phase-corrupted aperture illumination functions per

¹<https://casa.nrao.edu/casadocs/casa-5.6.0/introduction/casa6-installation-and-usage>

antenna to produce mis-pointed antenna primary beams (PBs), the code simply offsets airy disks to produce the latter. This allows for much greater angular resolution in the PBs and in turn fewer computational artifacts that would otherwise bias the image dynamic range and image fidelity measurements. This is possible because of the assumption of unblocked circular apertures with zero squint. For educational value, fully-functional code to perform the Fourier transform approach is included, commented-out, in the scripts.

- To significantly improve run-time, the code uses the `im.ft` tool for the visibility prediction step rather than `tclean`. This is the equivalent of `gridder='standard'` in `tclean`. It is not currently possible to run `im.ft` with `gridder='mosaic'`.
- Imaging uses `tclean` with `gridder='mosaic'` for the SBA simulations and `gridder='standard'` for the MA simulations. The former is technically inconsistent with the use of `im.ft` above, but testing indicates no detectable impact on the image corruption metrics examined in this memo. `tclean` has been hardcoded to assume natural weighting, though this could be easily modified if desired.
- The SBA code supports arbitrary true sky models. The MA code only supports true sky models that contain a single point source, so as to enable a simplification that significantly improves the accuracy of the `im.ft` step on long baselines (see MA script for details).
- The SBA simulations image a field of view that is larger than the PB. The MA simulations only image a relatively small fraction of the PB in a region about the injected source.
- `tclean` assumes `deconvolver='hogbom'` for single-channel simulations. Multi-channel simulations assume `deconvolver='mtmfs'` with recommended minimum `nterms=3`, the latter selected to minimize Taylor expansion errors as demonstrated in Figure 1.
- To prevent cleaning sidelobes or (pointing error induced) noise, clean masks are applied in `tclean` (specified with `usemask='user'` and a mask file; `auto-multithresh` is not used). These are defined with prior knowledge of the injected source locations. While this would not be possible for real observations, it is acceptable for the idealized imaging scenarios examined in this memo because it ensures negligible biasing of the image corruption metrics. The code also enables cleaning thresholds to be manually tuned to prevent excessive cleaning. To assist with data exploration, the code provides functionality to re-run a simulation with only the imaging functionality activated.

A challenge in constructing the simulation code was the critical need to control systematics arising from factors unrelated to pointing errors. In addition to the factors described above, care was taken in selecting the simple yet informative sky models described below in Section 4 to ensure that deconvolution errors (including potential over-cleaning biases) would not dominate any resulting image corruption metrics. Another challenge was the need to optimize compute time, memory footprint, and temporary disk space usage. The simulation scripts have been optimized (including parallelization) to ensure they are computationally tractable over these dimensions for the demonstration purposes of this memo. However, it is likely that even more efficient approaches will be required to successfully scale them to support simulations containing large numbers of channels and integrations.

3.1 Image Corruption Metrics

The simulations capture the following image corruption metrics:

1. Dynamic Range This is defined as

$$D_{observed} = \frac{\max(\text{Image})}{\text{rms}(\text{Residual})} \quad (3)$$

where the image and residual are flatnoise outputs from `tclean` (i.e. not PB corrected). For multi-channel simulations, `residual.tt0` is not the same as the residuals in `image.tt0` because the latter contains contributions from higher order residuals. To compensate, for wideband (multi-channel) simulations the code estimates the rms in `image.tt0` by measuring the median absolute deviation of all pixels with surface brightness < 0.1 Jy/beam and multiplies this by factor 1.48. This factor will recover the standard deviation for ideally normally distributed data. The standard deviation is not the same as rms, but should be approximately the same for these simulations with expected zero mean and median. (The rms isn't extracted because the few pixels near the 0.1 Jy/beam cutoff will otherwise bias the results.)

2. Image Fidelity This is defined for the SBA simulations as

$$F_{SBA} = 1 - \frac{|\sum \text{Image}_{PB,70\%} - \sum \text{Sky}_{True}|}{\sum \text{Sky}_{True}} \quad (4)$$

where `SkyTrue` is the true sky model injected into the simulation, and the sum over pixels in the PB-corrected image includes conversion from units of surface brightness to flux density and is only performed within the `pbthresh` $\geq 70\%$ mosaic PB contour. This contour level comfortably encloses the injected sources in the simulations for both the 3 and 7 pointing mosaics. The equation above is simplified for the MA simulations as

$$F_{MA} = 1 - \frac{|\max(\text{Image}_{PB}) - \sum \text{Sky}_{True}|}{\sum \text{Sky}_{True}} \quad (5)$$

where $\sum \text{Sky}_{True} = 1$ due to injection of a single point source with unit flux density. Note that alternate definitions of image fidelity are detailed by Mason (2019).

4 Results

4.1 Short Baseline Array

The SBA simulation code was run with combinations of the following parameters:

- Antenna configuration specified by SBA Revision C containing 19×6 m antennas with baselines ranging from 11 m to 60 m
- Spectral sampling centered in ngVLA band 4 at 27 GHz with two setups:
 - 1 channel
 - 13 equally-spaced channels (each with nominal 1 GHz width) with central frequencies spanning a wide fractional bandwidth between 21 GHz to 33 GHz
- Two hexagonal mosaics with pointing separations given by $\text{HPBW}/\sqrt{2}$ at 27 GHz:
 - 3 pointings arranged equilaterally about the field center
 - 7 pointings arranged with the field center coincident with the central pointing
- 1 min ‘integration’ time (recall lack of time average smearing in the simulation code); this also defines the pointing error coherence timescale
- Each pointing observed for 1 min with zero slew time, cycling through all pointings and repeating until the observation ends. The total observing time was set to 1 min per pointing, with the exception of a 3-pointing 1-channel simulation for which the total duration was set to 10 min per pointing.
- True sky model given by idealized extended source comprising a sum of 3 point sources located in a linear arrangement with maximum angular extent between injected source locations equal to 50% of the largest angular scale (LAS) of the SBA at 27 GHz. For reference, at 27 GHz the point spread function (PSF) full-width at half-maximum (FWHM) is $0.6'$, the LAS is $3.5'$, and the PB HPBW is $6.4'$. The multi-channel simulations were run with the spectral indices of all injected sources set to $\gamma = 0$ or $\gamma = -0.8$, following convention $S \propto \nu^\gamma$.
- Mean 2D radial pointing errors σ_2 ranging from $0''$ to $60''$. Simulations with $\sigma_2 = 0''$ were included to constrain image corruption limits arising purely from numerical noise.

The resulting image dynamic range and image fidelity metrics are displayed in Figure 2.

The measured image dynamic range limits are consistent with the values predicted using the equations presented in Section 2, including the \sqrt{M} scaling for the longer observations presented in the second row from the top. The predicted dynamic range limits for the 3- and 7-pointing mosaic simulations with 1 channel are identical (dotted curves in top row and third row from top, respectively). The consistency of the measured values with this curve in both panels indicates that even though the extended source is centered in the central pointing of the 7-pointing mosaic, the calculation for the 3-pointing mosaic described in Section 2 is appropriate for describing any mosaic in which common sources are detected over multiple pointings. The predicted dynamic range limits for the wideband 13-channel simulations (dotted curve in the lowest row) are also identical to those for the 3- and 7-pointing mosaic simulations with 1 channel. This is somewhat expected if multi-term deconvolution errors are negligible (as expected from Figure 1) and any spectral dependence arising from the combination of pointing errors and extended source emission averages out to negligible levels such that the overall behavior can be predicted using the central frequency alone. Note that a spectral contribu-

tion to the \sqrt{M} factor is not predicted because the primary beam response will be correlated between channels for any given antenna.

The image fidelity measurements indicate that pointing errors less than $\sim 30''$ ($\beta < 1/13$) result in minimal fidelity corruption ($F > 0.95$). The measurements appear to follow a functional form related to $F = 1 - f(\beta^2)$.

No significant difference is observed between measurements from simulations with different injected spectral indices. This is consistent with the interpretation above.

4.2 Main Array

The MA simulation code was run with combinations of the following parameters:

- Antenna configuration specified by MA Revision C containing 214×18 m antennas with baselines ranging from 30 m to 1000 km
- 1-channel spectral sampling centered in ngVLA band 4 at 27 GHz
- Single pointing observation
- 1 min ‘integration’ time (recall lack of time average smearing in the simulation code); this also defines the pointing error coherence timescale
- Simulations run with total observing time 1 min and 10 min
- True sky model given by two arrangements:
 - Single point source located at the field center
 - Single point source located at the half-power (-3 dB) radius
- Mean 2D radial pointing errors σ_2 ranging from $0''$ to $60''$. Simulations with $\sigma_2 = 0''$ were included to constrain image corruption limits arising purely from numerical noise.

The PSF for the MA simulations was found to be approximately 10x larger than would be expected from an array containing 1000 km baselines. This is a consequence of using natural weighting, as examined by Rosero (2019). For reference, this resulted in a PSF of ~ 20 mas at 27 GHz. Note that while the PSF is non-Gaussian with a significant plateau, deconvolved point sources were found to be clearly Gaussian with FWHM consistent with this PSF width. The PB HPBW is $2.1'$ at this frequency.

The resulting image dynamic range and image fidelity metrics are displayed in Figure 3.

The measured image dynamic range limits are consistent with the values predicted using the equations presented in Section 2, including the \sqrt{M} scaling for longer observations and reductions in dynamic range for the off-axis simulations, but only if the number of antennas in the array is assumed to be approximately half the 214 antennas used in the simulations. This is another expected consequence of using natural weighting, where for any given image resolution the bulk of the sensitivity from the MA will arise from approximately half the antennas in the array (Rosero 2019).

The image fidelity measurements indicate that pointing errors less than approximately $10''$ ($\beta < 1/13$) result in minimal fidelity corruption ($F > 0.95$). The measurements appear to follow a functional form related to $F = 1 - f(\beta^2)$.

Acknowledgements: I thank Urvashi Rau, Kumar Golap, and Brian Mason for helpful suggestions.

References

- Bhatnagar S., Cornwell T. J., Golap K., 2004, EVLA Memo No. 84
 Hales C. A., 2019, ngVLA Memo No. 60
 Kundert K., Rau U., Bergin E., Bhatnagar S., IEEE Trans. Ant. and Prop., 65, 2, 644
 Mason B., 2019, ngVLA Memo No. 67
 Perley, R. A. 1999, Synthesis Imaging in Radio Astronomy II, 180, 275
 Pety J., Gueth F., Guilloteau S., 2001, ALMA Memo No. 398
 Rosero V., 2019, ngVLA Memo No. 55
 Tsutsumi T., Morita K.-I., Hasegawa T., Pety J., 2004, ALMA Memo No. 488.1

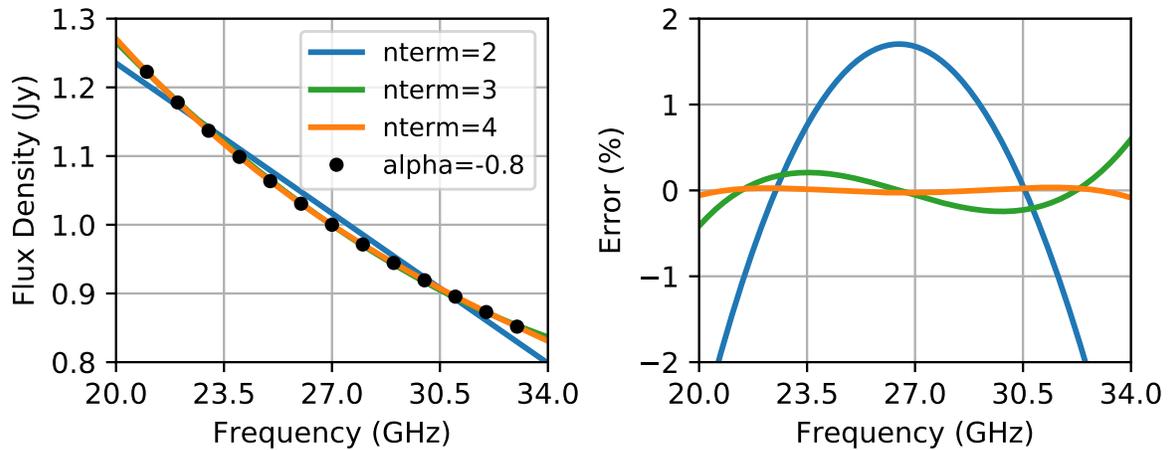


Figure 1 Flux density errors arising from different Taylor expansion orders when fitting a power law $S \propto \nu^\alpha$, relevant for `tclean`. Filled circles in the left panel indicate 13-channel sampling, each with idealized zero bandwidth, for an input spectrum with $\alpha = -0.8$. Curves indicate least-squares fits using multiple Taylor terms about 27 GHz. The error at 27 GHz is -18 dB, -33 dB, and -36 dB for 2, 3, and 4 Taylor terms, respectively.

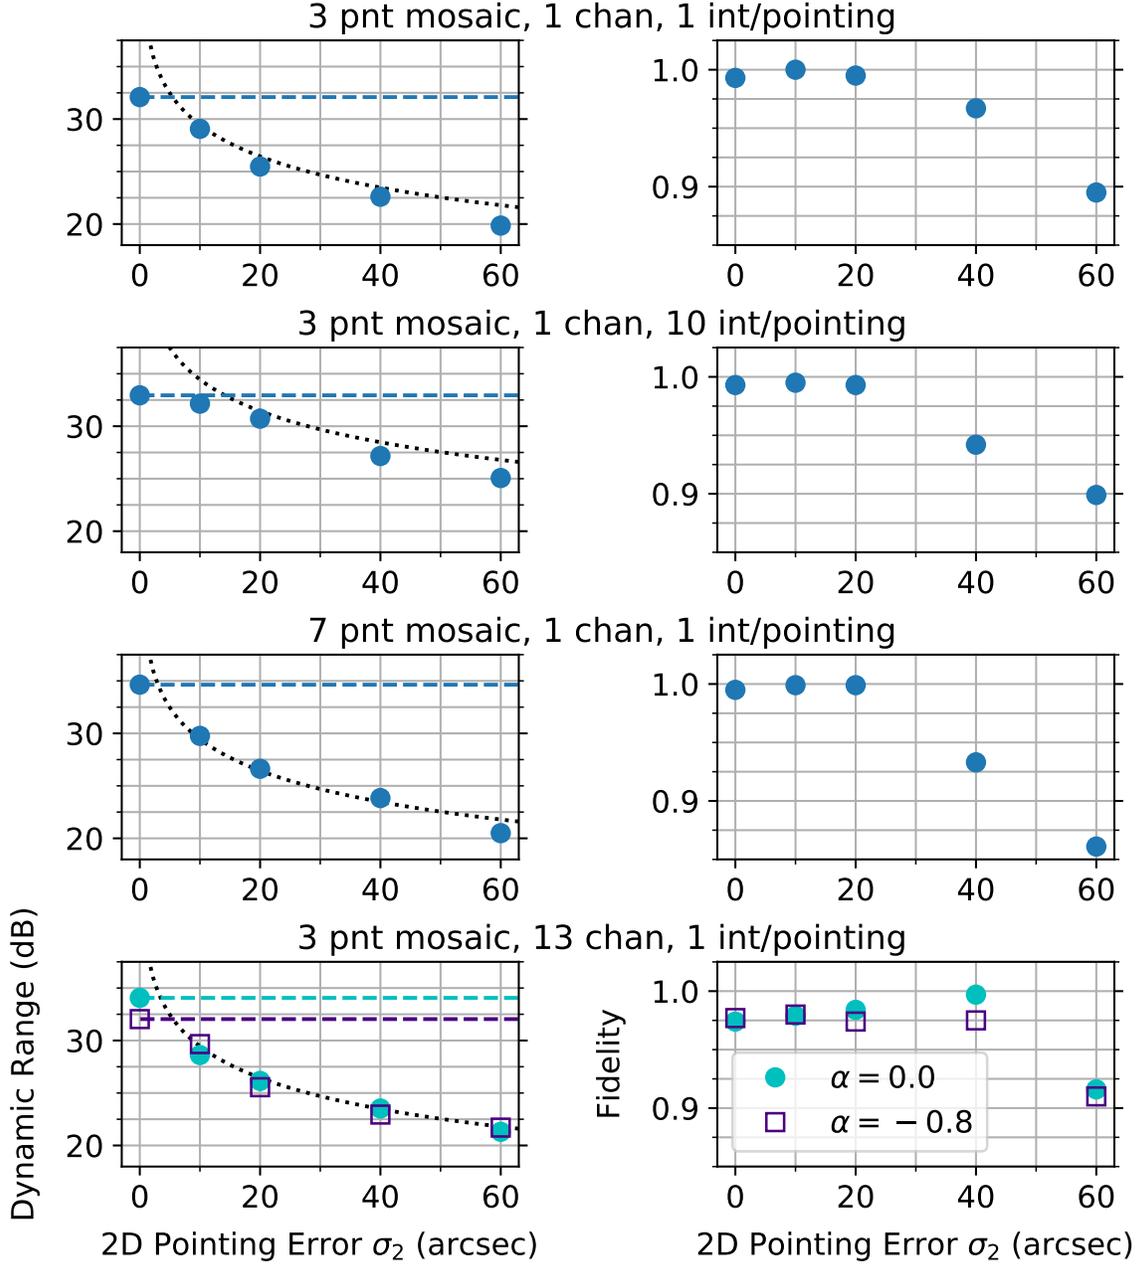


Figure 2 SBA simulation results for image dynamic range (left column) and image fidelity (right column) as functions of mean 2D radial pointing error σ_2 . The panels in each row present simulation results for the indicated number of pointings in the mosaic, spectral channels, and number of integrations per pointing throughout the observation. The horizontal dashed curves in the left panels pass through their respective $0''$ measurements and indicate the numerical limits of the simulations (i.e. other measurements in the panel cannot exceed this limit). The dotted curves in the left panels display the predicted image dynamic range limits arising from the equations presented in Section 2.

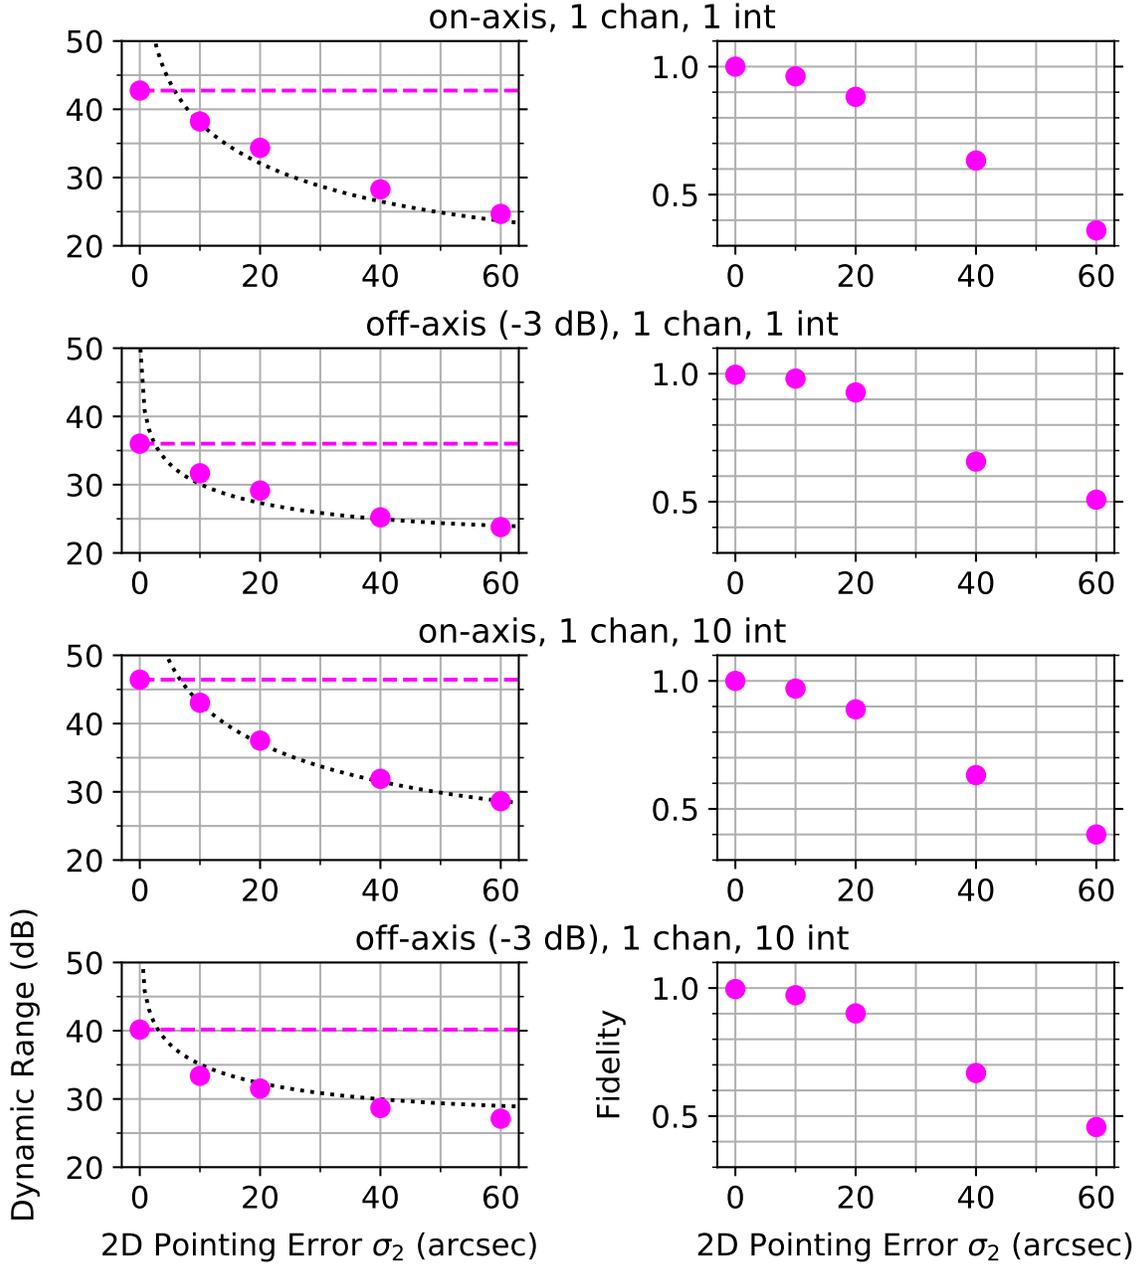


Figure 3 MA simulation results for image dynamic range (left column) and image fidelity (right column) as functions of mean 2D radial pointing error σ_2 . The panels in each row present simulation results for the indicated injected point source location, number of spectral channels, and number of integrations throughout the observation. The horizontal dashed curves in the left panels pass through their respective $0''$ measurements and indicate the numerical limits of the simulations (i.e. other measurements in the panel cannot exceed this limit). The dotted curves in the left panels display the predicted image dynamic range limits arising from the equations presented in Section 2, but assuming only half the number of antennas ($N = 107$) included in the simulations.