ALMA Memo 398

Impact of ACA on the Wide-Field Imaging Capabilities of ALMA

Jérôme PETY, Frédéric GUETH, Stéphane GUILLOTEAU

IRAM, Institut de Radio Astronomie Millimétrique
300 rue de la Piscine, F-38406 Saint Martin d'Hères

November 18, 2001

Abstract

This document presents the imaging simulations performed at IRAM to study the importance of the ALMA Compact Array (ACA) on the wide field imaging performances of ALMA. The simulations are based on a fast, approximate but sufficiently accurate method. The image reconstruction is based on two new CLEAN-based deconvolution methods for mosaics. The effects of *i)* pointing errors, *ii)* amplitude calibration errors and *iii)* atmospheric phase noise have been systematically studied. The effect of a poor knowledge of the primary beams has also been studied. Finally we investigated the impact of variations of the deconvolution parameters.

A large number of simulations have been performed, using several test images. This allowed us to disentangle artifacts due to particular random sampling of the errors from the genuine, average impact of these errors on all images. We find that adding ACA observations improves the image quality of ALMA in some types of images. More importantly, ACA provides robust results, when ALMA fails in presence of some type of errors, or on some images. One of the two new deconvolution techniques requires only a very moderate increase in computing time to include the ACA data.

Contents

1 Introduction

With the prospect of Japan joining the ALMA project in a 3-way partnership, it is necessary to explore what are the best scientific options for the possible extensions of the baseline ALMA proposal. These include a compact array of smaller antennas (the ALMA Compact Array – ACA), which would hopefully provide enhanced wide-field imaging capabilities by measuring the short spacings poorly or not sampled by ALMA.

In theory, all *uv* spacings can be measured by an homogeneous array like the ALMA baseline (64 12m–antennas), provided single-dish data is added and mosaicing is used. However, in practice, deconvolution methods do not necessarily recover the proper visibilities on intermediate baselines between 6 and 15 m. Moreover, pointing errors translate into phase errors proportional to the *uv* distance from the closest measured visibilities. Those phase errors are thus greatest around 9 m.

An array of small antennas, e.g. 7m–dishes, could measure directly these visibilities and be less sensitive to the pointing errors. Initial simulations (Pety et al. 2001 [5, 6], Yun 2001 [9], Morita 2001 [4]) showed that this is indeed the case. However, measured visibilities are also affected by other type of errors, like atmospheric phase noise and errors of amplitude calibration. No complete simulation including realistic values for these errors has yet been performed. Thus, it is not yet known whether the advantage provided by the ACA addition about the pointing errors remains significant in practical observing conditions.

This document presents in a comprehensive way the software developed at IRAM to do such simulations, and the results that have been obtained. In particular, it addresses the following problems:

- Can the short spacing information be recovered properly?
- Is the data obtained with ACA more immune to errors than data obtained with ALMA only?
- What happens when pointing errors induced by anomalous refraction are considered in a coherent way with the atmospheric phase errors?
- Are these results independent of the types of scientific problems?

Sections 2–4 presents the details of the methods used for the simulations. The readers interested in the results only may skip these sections and proceed directly to Sections 5–8, or even to the conclusions (Section 9).

2 Simulator Overview

Our goal is to compare wide–field imaging capabilities of ALMA with or without the use of ACA data. This section gives an overview of the steps taken to do this. In short, we need to simulate *i)* interferometric observations with ALMA and ACA and *ii)* single-dish observations with a limited number of 12m-antennas. To make these simulations realistic, different type of errors are introduced (a thorough description is given in Section 3). The different simulated observations must then be merged and deconvolved to produce the simulated images (cf. Section 4). A comparison between those simulated images and the input model finally allows us to evaluate the contribution of ACA to the image quality.

2.1 The Simulation Process

The simulator includes the following steps:

Interferometer simulation

The same routines are used to simulate ALMA or ACA observations (only the input parameters differ). For simplicity sake, simulation is made in 3 steps: amplitude and phase errors are first encapsulated as complex gains; ideal visibilities are then computed from the input model image; error gains are finally applied to the ideal visibilities.

1. *uv* coverage
   The *uv* coverage is generated from the array configuration, source declination, and observing time. A short timescale *uv* table corresponding to the baselines sampled every second is generated. It initially contains perfect gains (Amplitude 1, Phase 0).
2. Instrumental gain.
   Gains are then corrupted by the phase errors due to the atmospheric phase screen, and by the amplitude errors. For convenience, effects resulting from phase calibration are included in this step, i.e. the final gains contained in the *uv* coverage table are **calibrated** gains. A table of pointing errors induced by anomalous refraction is also computed at this step.
3. Time average.
   The gain table is then time-averaged according to the selected integration time (dictated either by the Nyquist rate, or the observing sequence). Correction for decorrelation is applied at this step.
4. Model visibilities.
   Generation of the model visibilities, including pointing errors, is made only on the integration–time basis. Hence, the pointing errors tables (induced by anomalous refraction and/or from a model including thermal or wind effect) are first time averaged. Visibilities are then computed for each integration period using the effective pointing centers for each antenna. Pointing errors induced by anomalous refraction can be taken into account.

   Mosaics are simulated by adding the field offsets to the pointing error at the beginning of this step. The output *uv* table is sorted to produce one *uv* table per field. Note that the procedure assumes that the fields are observed in a loop (i.e. in a snapshot mode), each visibility dump corresponding to another position in the sky. The total integration time is truncated so that an integer number of cycles is observed.
5. Corrupted visibilities.
   The visibilities computed in the previous step are multiplied by the gain table obtained in step 3 to account for phase and amplitude errors.

Single-Dish simulation

1. Single-Dish observations.
   The model image is convolved by the antenna lobe and the intensity is then estimated at *(i)* the position of the mosaic fields (for merging with ALMA visibilities), and *(ii)* on an extended grid (for merging with ACA visibilities), with typically 3 pixels per beam. If the required position does not coincide with a pixel center, a bilinear interpolation from the neighbor pixels is performed.

Pointing errors can be included by estimating the intensity at the position corrupted by the pointing offset. Thermal noise can then be added to the data and the corresponding weight is stored. Finally, an error of amplitude calibration can be simulated.

2. Extracting Single-Dish pseudo-visibilities.
   If the Single-Dish data are to be merged with the ACA data, then "pseudo-visibilities" are computed from the Single-Dish map (see Section 4.2.2 for a detailed description of the algorithm).

Imaging

1. Merging interferometric and Single-Dish visibilities.
   The merging of the two datasets is done for each field. The relative weight between the two datasets is an adjustable parameter (see Section 4.4.1).
2. Mosaic reconstruction.
   Dirty images are produced for each mosaic field, and are then linearly combined to produce the dirty mosaic (cf. Section 4.1).
3. Deconvolution.
   The mosaics are deconvolved with the appropriate algorithm (see Section 4). Depending on the initial setup of the simulator, six images can possibly be produced: ALMA-only, ACA-only, ALMA+SD, ACA+SD, ALMA+ACA, ALMA & ACA+SD.
4. Hybridization.
   If required, two images can be combined in the *uv* plane to produce the **hybrid** deconvolved image (see Section 4.3.2). This is done in practice with ACA+SD and ALMA+SD.

Comparison between model and simulation

For each of the 7 possible simulated images, a comparison is performed with the original image, convolved with the appropriate CLEAN beam. A **fidelity** (see Section 5) image is computed. The comparison is performed both in the image and the *uv* plane.

2.2 The Simulator Interface

In order to perform the various steps involved in the simulation, flexible procedures as well as a window interface were written. The top part of Fig. 1 shows the control window of the simulator. The user has to define 3 global parameters: the input model image, the output directory, and the type of simulation. 7 different types of simulations can be made: ALMA only, ACA only, ALMA+ACA, ALMA+SD, ACA+SD, **all** (= ALMA & ACA+SD), and **hybrid** (*uv* plane hybridization). All results are stored in the same output directory, so that a simulation requiring results from another one shares the same files when needed.

A simulation actually consists in 3 different steps, each with its own associated button: `COMPUTE` the observation results, `COMPARE` with the original data, and `DISPLAY` the comparison. The `COMPUTE` step can be replaced by a `LOAD` step if results from a previous simulation exist. In addition an `EXPERT` button allows access to individual steps in the computations (not recommended, but useful for debugging new features).

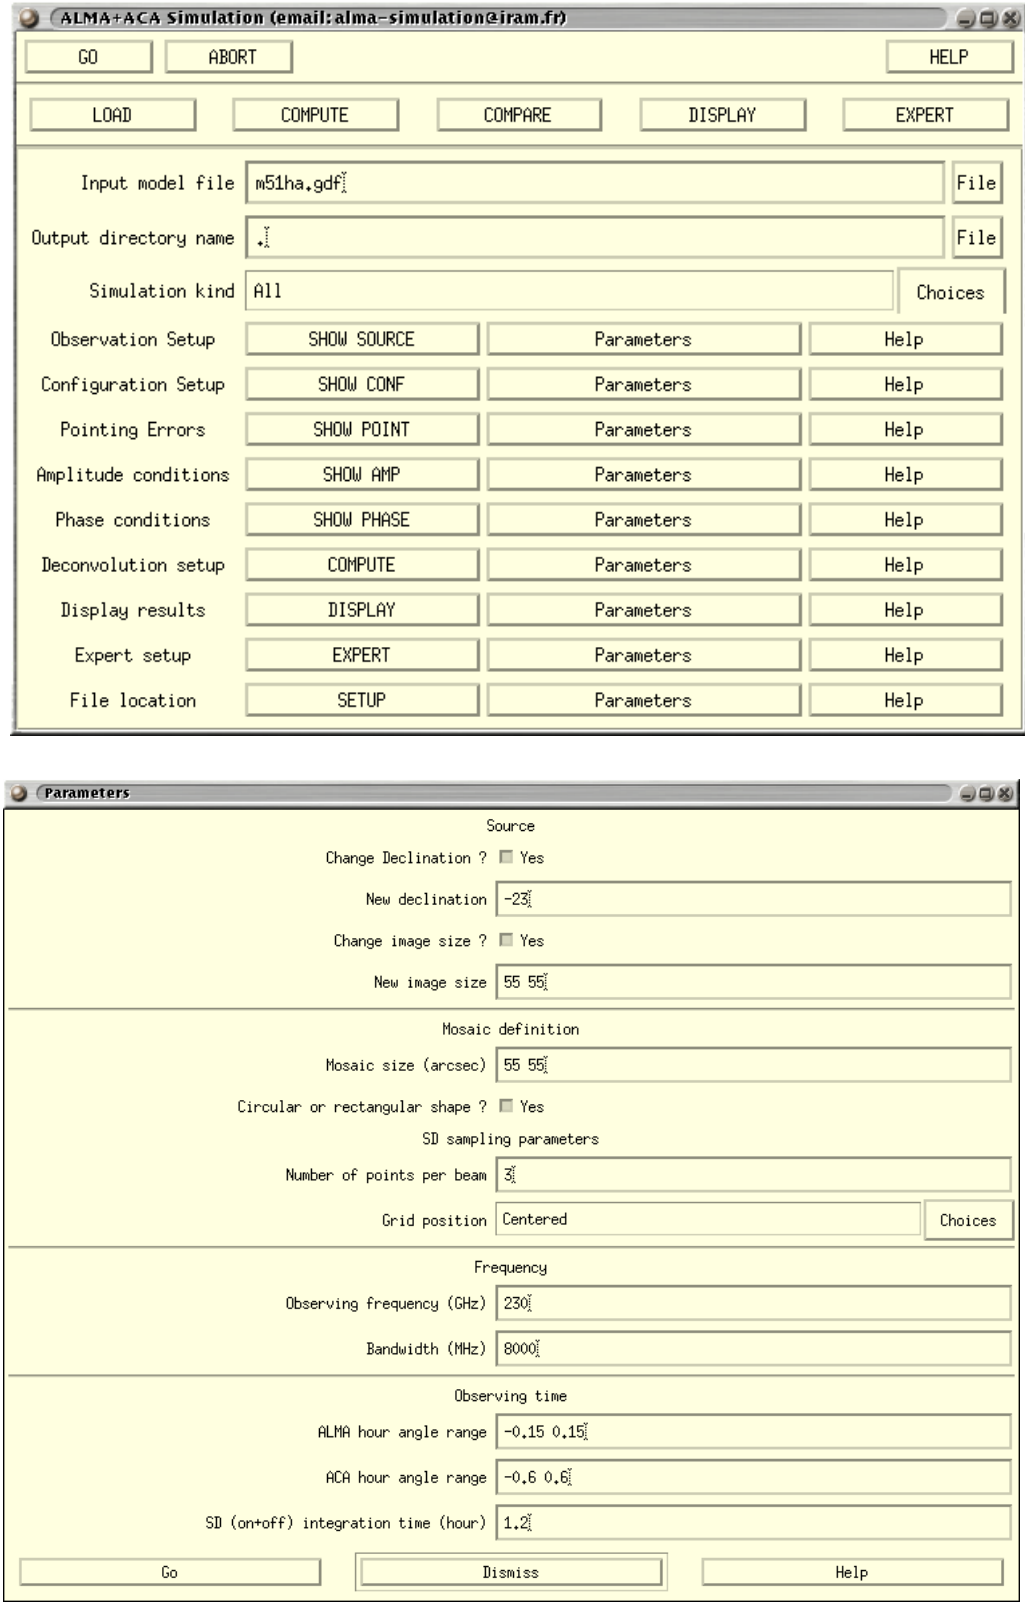

Figure 1: Top: Control window of the simulation package. Bottom: Window obtained by clicking on the "parameter" button in the main window, allowing one to modify the default parameters (here, of the `Observation Setup`).

The control window includes definition buttons, which allow one to setup simulation parameters (through pop-up windows as shown on the bottom part of Fig. 1):

`SHOW SOURCE` defines the source related observation parameters: source declination, image size, mosaic size and shape, parameters of the single-dish map, observing frequency and bandwidth, hour angle range for ALMA and ACA, and single-dish integration time.

`SHOW CONF` defines the array configuration parameters, of ALMA (style and size of the array configuration), ACA (size of antennas, style and size of the array configuration), and Single-Dish (number and size of antennas used for single dish).

`SHOW POINT` defines the pointing errors. For each sub-array (ALMA, ACA or SD), it specifies the error model, the rms value of the error, and a filename in case of table-driven errors.

`SHOW AMP` defines the amplitude conditions: errors and noise. Noise is derived from specified receiver temperature and zenith opacity. Amplitude errors are specified by the rms value, drift value and drift time constant.

`SHOW PHASE` defines the atmospheric phase conditions. They are specified by the phase screens for ALMA and ACA, the phase noise at 300 m, wind speed and direction, the screen altitude, the calibrator position, the integration time on source and calibrator, the parameters of the Water Vapor Radiometry, and the calibration method. In addition, a parameter indicates whether anomalous refraction should be computed and added to the pointing errors.

`COMPUTE` activates the simulation. This button has a number of parameters to specify the details of imaging and deconvolution: the Single-Dish weight, the primary beam truncation level for the mosaic, the number of CLEAN components, the CLEAN loop gain. When hybridizing different simulated images in the *uv* plane, the images and the *uv* transition radius to be used can also be selected here.

`DISPLAY` displays the results. Several displays are possible:

- The fields position;
- The single-dish image;
- The *uv* weights;
- The synthesized beam(s);
- The flux convergence, i.e. the cumulative flux as function of CLEAN component number;
- The image plane comparison: the model, simulated, difference, and fidelity images are plotted; it also presents the cumulative histogram of the fidelity;
- The *uv* plane comparison: the visibility amplitudes are plotted for the model and the difference; the fidelity directly computed in the *uv* plane as well as a radial average of the fidelity are also shown;
- Two different types of summary plots presenting a combination of the previous ones;
- The phase calibration results when available.

`EXPERT` toggles some controls to allow "experts" to force performing some simulation steps. This should only be used for debugging or special simulations.

`SETUP` specifies the location of the model images, pointing error tables, and phase screens.

The input parameters of the simulation and the intermediate results are stored in files, thus allowing one to re-run part of the simulation using previously computed data, i.e. without starting from the very beginning. Only the incomplete simulation steps are re-done when needed.

2.3 The Simulation Pipeline and Archive

On top of the GILDAS simulator, a pipeline was written to run automatically all the procedures and save the results. This allows us to run a large number of simulations to assess the impact of the different input parameters (e.g. short spacings addition, noise level, pointing errors, etc.) on the image quality.

The pipeline maintains a database of results, which can be browsed through GILDAS procedures to select and display part of it. More than 2000 different simulations were performed in the course of the ACA study, leading to more than 100 GBytes of results.

2.4 Environment of Development and Performances

The simulation package was developed and implemented in the GILDAS environment, used for the data reduction softwares of the IRAM instruments. As the whole GILDAS environment, the simulator runs under HPUX, Linux (RedHat), or MS-Windows. The simulator uses the MAPPING subpart of GILDAS. Starting from the existing tools, a number of new algorithms were developed and implemented, as e.g. the simulation of single-dish measurements, pointing errors models, the estimation of visibilities in presence of pointing errors, or the deconvolution of data taken with an heterogeneous array. In addition, several existing tasks were improved to be more flexible, and/or to save some memory space, and/or to speed up the computation time. For that purpose, part of the MAPPING software code was translated from FORTRAN-77 to FORTRAN-90, which allows an easier implementation of mathematical algorithms.

The amount of time needed to perform a simulation depends on the size of the problem. On a 1 GHz AMD Athlon PC with 756 Mbytes of memory, the simulation of a 7–fields mosaic observed with ALMA takes only a few minutes. The simulation of ALMA & ACA+SD is somewhat longer, $\sim$15–20 minutes, because the deconvolution of an heterogeneous array is rather slow. As a rule, the imaging part (Fourier Transform and deconvolution) is always the time-consuming step of the simulation.

3 Toward Realistic Simulations: Including Errors

The simulation engine has been designed to handle pointing, phase and amplitude errors. Speed was the major design goal. Accordingly, the simulations have limited accuracy; in the absence of error, the relative precision of the simulation is about 1 part in a thousand. However, as shown below, the inclusion of typical error values completely dominates the image quality. The relative inaccuracy of our modeling is thus not a problem.

3.1 Pointing Errors

3.1.1 Error Models

Pointing errors include:

- Systematic offset: a global mispointing of the mosaic.
- Random offsets per antenna: such offsets can result for example from measurement errors in the last pointing calibration.
- Wind induced effects: a randomly variable component, with a range of direction and magnitude, and a correlation coefficient between antennas. The timescale of variation is of order 1 second.
- Thermally induced effects: a drift with time, with slowly varying direction and a large correlation coefficient between antennas. This large correlation is used to reflect the common deformation mode of the antennas due to solar illumination. The timescale for variation is of order 10 to 30 minutes.
- The so-called "*anomalous*" or rather *Dynamic* refraction. The *anomalous* refraction is due to time variations of the refractive index of the wet air along the line of sight. This is a perfectly "normal" and unavoidable process, so we prefer to call it *Dynamic* refration, by opposition to the *Static* refraction due to the static part of the atmosphere.

The derivation of dynamic refraction is described in Section 3.2.3. The other terms are computed for each antenna on a 1 second timescale basis, and stored in a pointing error table. Time averaging is then applied to derive the mean pointing error on the integration timescale. Several tables of pointing errors have been generated. The use of a pre-calculated pointing error table allows us to apply the same errors to the ALMA and ALMA+ACA simulations, and hence compare intrinsic differences rather than effects due to different generation of random numbers.

3.1.2 Simulations of the Corrupted Visibilities

For a single field, the observed visibility is the Fourier Transform of the multiplication of the sky brightness distribution by the primary beam of the antenna (i.e. the convolution of the Fourier Transform of the sky brightness by the antenna transfer function, which is simply the Fourier Transform of the antenna beam). For rapidly varying pointing errors, as this is the case in our simulations, this quantity has to be computed for each integration time or dump. Rather than computing this convolution by a multiplication in the image plane followed by Fourier transform (which would be needed for each integration time), we perform it directly in the *uv* plane. The visibilities are estimated by the convolution of the (gridded) FFT of the sky brightness distribution by an analytic form of the antenna transfer function. We use a truncated Gaussian illumination, with 11.5 dB edge taper[1]. Pointing errors are simply handled by applying the appropriate phase term to the antenna illumination pattern. Furthermore, the convolution is approximated by a weighted sum of the neighbor points. Since the illumination pattern has an outer radius equal to the antenna diameter, this is a local operation in the *uv* plane, to be performed at each required *uv* position.

[1] The transfer function is the autocorrelation of the antenna illumination.

3.2 Phase Errors

Phase errors are directly proportional to pathlength errors. Phase errors can thus be introduced into the simulation, using a pre-calculated table of pathlength errors. Only pathlength errors due to the atmosphere are used in this memo, although introduction of other types of pathlength errors would be possible (e.g. baseline errors).

3.2.1 Atmospheric Phase Screen

To compute the atmospheric phase errors, we generate a 2-D phase screen on a grid with sufficient spatial resolution to sample the antenna diameter (typically 4-m pixels for ALMA, and 2-m pixels for ACA). The phase screen is generated in the Fourier plane with the constraint that its 2nd order structure function is a combination of 3 power laws in 3 different spatial ranges. The resulting phase screen is averaged over the effective dish diameter. Because of the size limitation imposed by the FFT, very long phase screen are built as a linear combination of independently generated screens. This is correct since atmospheric pathlength variations are completely uncorrelated at large distances.

Observations through the phase screen are simulated in a realistic way. The observing direction through the screen is computed every second for the source as well as for the calibrator. Moreover, the phase screen is assumed to move with a given wind speed along one direction. For simplicity sake, the wind direction is assumed to be constant in time. This simplification should only have a small effect on the phase error since the phase screen is isotropic (see however the dynamic refraction paragraph). The free parameters are: the rms phase on the 300-m baselines, the angular distance between source and calibrator, the direction of the wind with respect to the source-calibrator vector, and the wind speed and altitude of the phase screen above the ground.

The use of the 2-D approximation prohibits exact handling of the elevation dependence of the phase function. We instead used a simple prescription in which the phase error scales with the square root of the elevation. An example of a phase screen and its gradients is given in Fig. 2.

3.2.2 Phase Calibration

Since source and calibrator are observed through the same phase screen, standard phase calibration techniques can also be applied to the simulated data. 4 different calibration methods have been implemented:

1. Standard calibration: a continuous curve (polynomial or smoothing spline) is fitted through the time sequence of calibrator phases , and removed from the source phases.

2. Fast switching: a linear interpolation of the phase is performed between two consecutive calibrator measurements, and removed from the source phases.

3. Water Vapor Radiometry: a WVR might be used to correct the atmospheric phases in real time. We modeled the corrected phase as

$$P_{\rm WVR}(t) = P(0) + \chi(P(t) - P(0)) + N(t) \qquad (1)$$

where $P(t)$ is the original atmospheric phase, and $t = 0$ is the beginning of the observing sequence. χ is the precision of the radiometer: $\chi = 0$ represents a perfect correction, $\chi = 1$ no correction. $N(t)$ is the phase noise introduced by the WVR. In our model, we assume the WVR is not able to follow the atmospheric phase when slewing from the calibrator to the source ("*non-connecting*" WVR), and thus reset

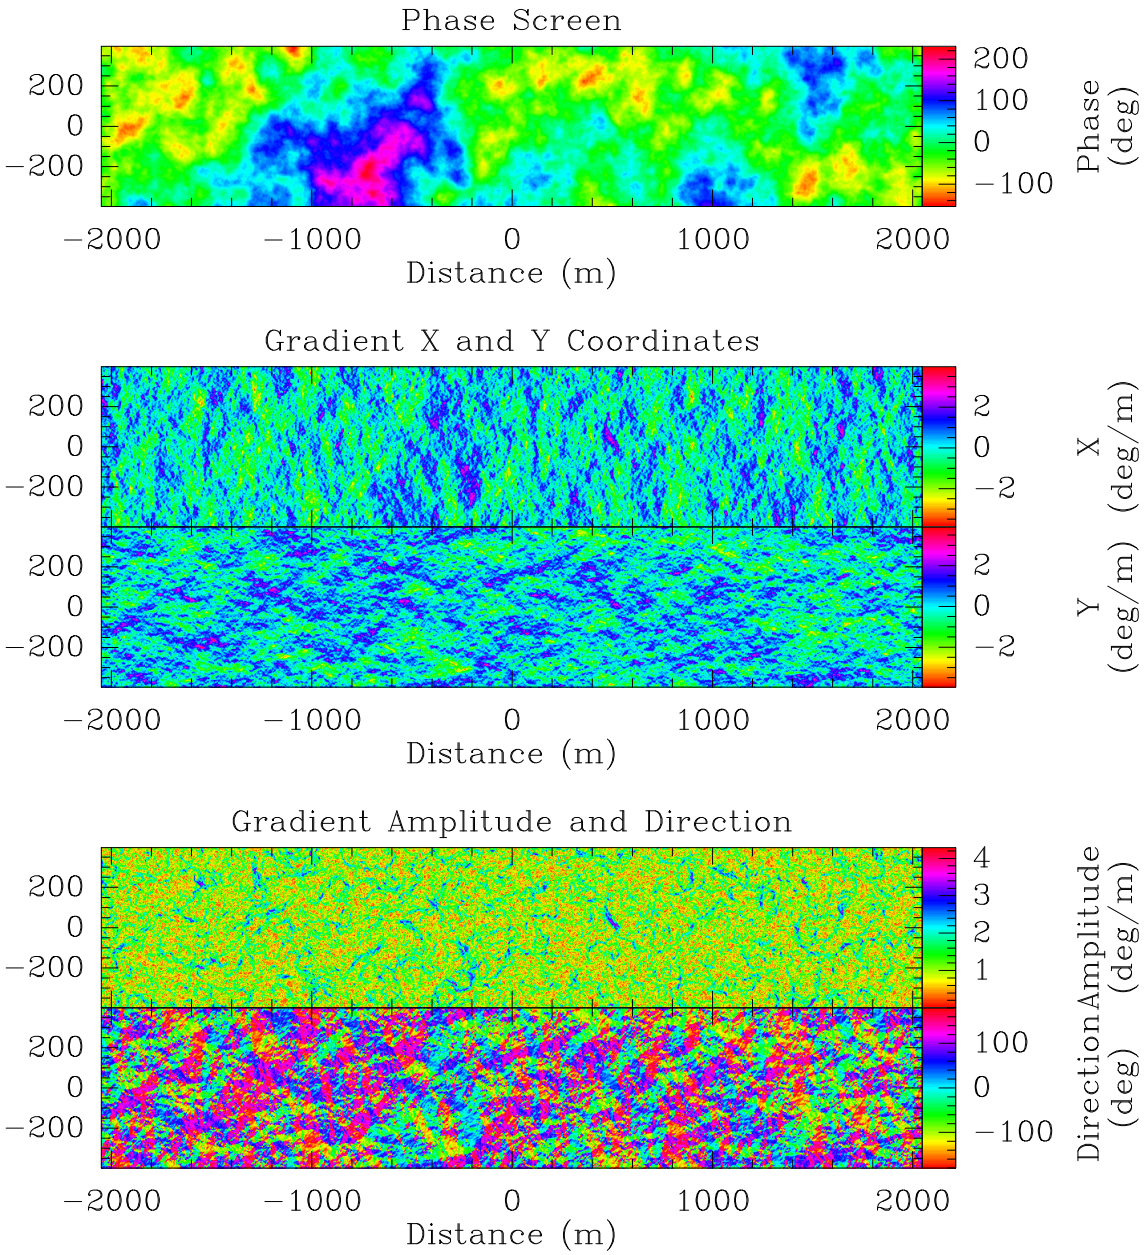

Figure 2: Example of a phase screen with its associated phase gradients, also shown as magnitude and direction.

$t = 0$ at each change between source and calibrator. A subsequent phase calibration is thus required.

4. Combined WVR and Fast Switching: in this mode, the WVR is used to correct the short term phase fluctuations, and the phase on the last calibration $P_{\rm cal}$ is subtracted from $P_{\rm WVR}(t)$. This is an optimum scheme for "*non-connecting*" WVR.

Fig. 3 shows the phase structure function before and after phase calibration for Option 2 and 4. Option 4 gives the best result on baselines longer than about 80-m. However, there still is a residual phase error after calibration even with a perfect WVR ($\chi = 0$ and $N(t) = 0$). This is due to the angular distance between the source and calibrator, and the limited antenna speed which imposes a time delay between the source and calibrator

measurements. Only a *connecting* WVR, which allow (partial) pathlength correction even when slewing from source to calibrator, would possibly give better results. Option 4 gives better result than Option 3 (not shown here) because the time delay is half the switching period in Option 3, while it is only the slew time in Option 4.

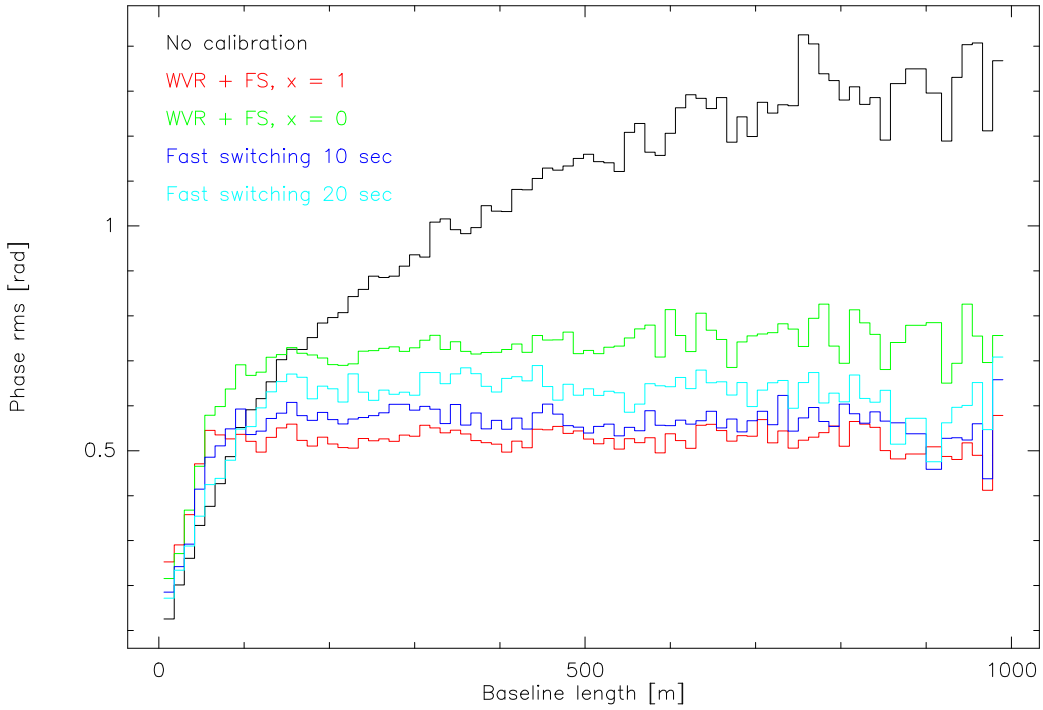

Figure 3: Residual rms phase errors as function of baseline length after application of different calibration methods. The calibrator was 2 degrees away from the source, and the phase screen was 1000-m above the site, moving at a wind speed of 10 m/s.

Note also that calibration **increases** the phase noise on short baselines. Moreover, the shorter the calibration timescale, the larger the phase noise on short baselines increases. This indicates that a better calibration strategy for ACA, whose baselines always are shorter than 60 m, would be to use Option 3 with a long calibration timescale.

3.2.3 Dynamic Refraction

Dynamic (anomalous) refraction is directly proportional to the phase gradient. We thus compute the phase gradient associated to the phase screen, in order to obtain a coherent derivation of the dynamic refraction term. As the phase screen, the phase gradient is averaged over the effective dish diameter.

In practice, we did not compute the effect of the resulting pointing error every second, and average the visibilities afterward, but rather compute the mean pointing error on each integration time. This procedure is justified because the theoretical (error free) visibilities are not changing significantly within each integration time (since we sample better than the Nyquist rate).

Computing the mean pointing error in each integration time is equivalent to averaging the phase gradients in an anisotropic way, i.e. convolving them with a function elongated along the wind direction. It can be seen from Fig. 2 that after such an anisotropic smooth-

ing, one of the (smoothed) gradient components will be smaller than the other, since each gradient has a different preferential direction . The resulting gradient is larger across than along the wind direction, implying that dynamic (anomalous) refraction has a privileged direction, related to the wind. As surprising as it looks, this is an intrinsic property of the gradients of turbulent fields.

In actual conditions, this anisotropy may be small because the wind direction has some random variations. Moreover, it would be difficult to observe because other pointing errors will also have preferential directions with respect to the wind (such as the errors induced by wind pressure on the antenna).

3.3 Amplitude Errors

Amplitude errors mostly originate from *i)* uncertainties and variations in the atmospheric transmission, and *ii)* slow antenna gain variations due to dish distortion and defocus, especially at high frequencies. Receiver gain fluctuations and drifts should remain at a much lower level than the two previous, unavoidable, effects.

3.3.1 Interferometer Model

The error on the amplitude gain is modeled as the sum of an offset and a drift with time. The offset and drift values are randomly reset at each calibration for each antennas. The typical amplitude calibration timescale has been set to 15 minutes. So far, no correlation coefficient between the amplitude errors of the various antennas has been used, although some of the effects will have a common mode (e.g. it is likely that thermal deformation of the dishes are similar in all antennas).

It is assumed that the amplitude drop due to decorrelation (due to the phase fluctuations during the integration time) is completely calibrated out. Failure to do so results in systematic under-scalings of the visibilities, but has no significant influence on the morphology, since the phase noise is essentially constant as function of baseline length after calibration (see Fig. 3).

3.3.2 Single-Dish Model

As for the interferometer, the error on the amplitude gain is modeled as the sum of an offset and a drift with time, whose values are randomly reset every 15 minutes for each antenna. In addition, an offset common to all the antennas can also be added to the amplitude gain.

For the single-dish data, the effects of errors on flux densities highly depends on the observing strategy. How many antennas are used, do they scan the source or do they point on a few observing direction, what are the possible scanning directions, what averaging is performed become a required knowledge to derive the final resulting error on the single-dish image. In the absence of a prescription for the Single-Dish observing strategy, we implemented the two simplest approaches depending on the kind of interferometric data that Single-Dish will complement.

Combination with ACA data A single-dish map is here useful to provide the short spacing information (see Section 4.2.1). We thus assume that a single-dish map is obtained for each antenna by scanning the appropriate region as slowly as allowed by the available observing time. At each second, the 4 used antennas look at the same direction. The 4 maps obtained with the 4 antennas are later averaged together. For example, for a 7 field mosaic, each single-dish map is made of 21×21 points sampled at $1/3^{\rm rd}$ of the beamwidth. Obviously, obtaining the maps by scanning as fast as possible, and crossing the scanning

direction between antennas at each second would result in more homogeneous maps. Hence, the selected method somewhat maximizes the distortion one may expect from amplitude errors.

Combination with ALMA data When combining with ALMA data, only the total flux for each field can be used (see Section 4.2.2). We thus assume that each field center is observed at least once per independent amplitude calibration sample (e.g. every 15 minutes in our typical simulation). For example, for a 7 field mosaic, only 7 directions on the sky are observed. Compared to the previous map, the amplitude error is thus smaller by a factor of $\sim \sqrt{21 \times 21/7} \sim 8$ in the typical simulation. Care is also taken that exactly the same amplitude errors are applied to these single-dish fluxes than to the single-dish map. Hence, for the 7 field mosaic, each of the 7 pointing directions is observed 63 times to use the same 441 error values as used for the single-dish map required for ACA. Further reduction of the amplitude error will result from the averaging over the 4 Single-Dish antennas, because of the uncorrelated part of the amplitude errors.

Comparison of the two observing strategies The simulation thus somewhat disadvantages the ACA observations (which rely on a Single-Dish map) with respect to ALMA observations. The opposite is perhaps more likely to happen in reality. Indeed, assuming that we can actually directly measure the flux density for each field center requires an observing method such as wobbler switching to a blank field. In many circumstances, such blank fields will not exist, and Single-Dish flux densities for the mosaic fields may have to be derived by interpolation from a fully sampled On–The–Fly map, rather than from direct observations of the field centers. Even with full sampling, the interpolation errors could easily be of order 2-3 %.

3.4 Thermal Noise

Thermal noise can be accounted for in our simulation. The thermal noise is derived from $T_{\rm sys}$ (taking into account the variation with elevation), the integration time, and the bandwidth. This is done for Single-Dish data as well as for interferometric observations. The simulation process also allows us to add random phase noise to the visibilities. Such phase noise is expected, especially when using a WVR.

4 Mosaic Deconvolution Techniques

Deconvolution of the raw images obtained with ALMA and ACA is required to provide correct comparison between the initial image and the simulation results. We used CLEAN-based deconvolution methods, which we summarize in this section.

4.1 Data from an Homogeneous Array (ALMA-only or ACA-only)

The deconvolution of a mosaic observed from an homogeneous interferometer (i.e. all array antennas have the same diameter) is done using a CLEAN-based method developed for the IRAM Plateau de Bure interferometer (see e.g. the chapter on mosaicing in the proceedings of IRAM millimeter interferometry summer school). In short: the dirty mosaic is reconstructed as a linear combination of all dirty maps,

$$J = \left(\sum_i \frac{B_i}{\sigma_i^2} F_i \right) \Big/ \left(\sum_i \frac{B_i^2}{\sigma_i^2} \right) , \qquad (2)$$

where F_i, B_i and σ_i respectively are the dirty map, the primary beam and the rms noise of each field. The mosaic is thus homogeneous to the sky brightness distribution, but the noise level is not uniform and strongly increases at the edges of the field of view. As a consequence, using the classical CLEAN method can be dangerous, as noise peaks might be selected as CLEAN components. The method used for the Plateau de Bure data consists in selecting the CLEAN components on a modified mosaic: the initial mosaic is normalized by the noise level distribution. Hence, this modified mosaic is homogeneous to the signal-to-noise ratio.

Truncated primary beams B_i are used instead of the theoretical beam to avoid propagation of errors at beam edges. The truncation threshold is one free parameter in the method. The algorithm turns out to give excellent results, both in terms of image quality and computation time. It has been used during the last few years to deconvolve the mosaics observed with the Plateau de Bure interferometer.

4.2 Using Single-Dish Data

In the presence of extended structure, the previous algorithm cannot recover all structures[2]. Inclusion of the short spacing information is required.

4.2.1 ALMA and Single-Dish

In this case of an homogeneous array, the short spacings are measured with the same antennas as the interferometric data. As the measured visibilities results from the convolution of the Fourier Transform of the model with the antenna transfer function (i.e. a local average, cf. Section 3.1.2), single dish measurements can only give the *zero spacing* in this case. This is done by inserting for each field the *zero spacing flux* (as measured in single-dish mode) prior to the imaging. The mosaic deconvolution then proceeds as before. The method provides good convergence properties.

4.2.2 ACA and Single-Dish

When complementing ACA interferometric data, the short spacings will come from the ALMA antennas. It is then possible to take advantage of the larger diameter (D) of the ALMA antennas to derive short spacing informations for the small antennas of diameter d up to distance $D - d$ in the uv plane (where the convolution to compute the visibility makes sense). This is done by the following operations:

1. The Single-Dish measurements are re-gridded and then FFTed into the uv domain.
2. The data are corrected for the single-dish beam by division by its Fourier Transform (truncated to the antenna diameter).
3. The data are FFTed back to the image plane where the ACA primary beam is applied (by multiplication).
4. The visibilities are then interpolated from a regular grid after yet another FFT.
5. In the case of a mosaic, the two last operations are performed for each pointing center.

Visibilities up to $D - d$ are finally extracted and merged with the ACA data before proceeding to the mosaic stage as in Section 4.1.

[2] In theory, with an infinite size mosaic, short spacings can be recovered, but the zero spacing is recovered only with zero weight...

4.3 Data from an Heterogeneous Array (ALMA and ACA together)

When using data from the ALMA and ACA arrays, the imaging problem becomes more complex, since two different types of primary beams now exist. We have developed two imaging methods with can handle such situation, provided no correlations between small and large antennas are performed. In the first method, the deconvolution operates at the same time on both types of data. In the second method, ALMA and ACA mosaics are separately deconvolved and then combined in the *uv* plane. This second method has a very small computation overhead compared to the deconvolution of the ALMA mosaic alone.

4.3.1 Joint Deconvolution

Description Starting from two *uv* tables corresponding to the observations of the same source with two arrays, the following operations are performed:

1. Fourier Transform of the data, and reconstruction of two separate mosaics (the ALMA and ACA ones). Although the field positions do not need to coincide in the observations performed by the two arrays, this is the case with the current version of the simulator.

2. Major cycles are performed, in which CLEAN components are identified in the residual image with the highest signal-to-noise ratio. A very satisfying method (in terms of image quality and convergence speed) is to use the Clark method [1] if the components are to be found on the ALMA image, and the Steer-Dewdney-Ito approach [8] if they are to be found on the ACA image.

3. The components found in each major cycle are subtracted from both residual images. If they were identified on the ALMA image, a compression allows us to re-sample them on the ACA grid. However, the components found on the ACA image need a more complex processing to make sure that the same quantity is subtracted from both residual images (otherwise the method would diverge). We first derive from each ACA component a set of contiguous positions on the ALMA grid, such that they cover exactly the area of the pixel of the ACA grid; these components are then subtracted from the ALMA residual image; they are also compressed (via FFT) back to the ACA grid, to be subtracted from the ACA residual image. This back-and-forth method allows us to get rid of the ambiguities (interpolation errors, aliasing) that would occur if the ACA components would have been directly interpolated on the ALMA grid.

4. Finally, all CLEAN components identified by the algorithm are convolved by the ALMA clean beam (highest angular resolution), and the weighted residual images are added.

Short spacing information can also be introduced in this method, either in the ALMA data as before (see Section 4.2.1), or (more logically) in the ACA data (Section 4.2.2), or even in both. In practice, we incorporated the short spacings in the ACA data. As with the more simple homogeneous problems, the relative weight of Single-Dish data compared to interferometric data is a key parameter in the convergence speed. A poorly adjusted weight can slow down considerably the deconvolution. In this work, we always used the weights imposed by the noise level for this purpose (see Section 4.4.1).

Convergence Properties The convergence properties of this method are more complex (and less satisfying) than those of standard CLEAN. The algorithm usually quickly recovers the total flux, but then often needs a large number of components to locate the fine structures. Flux overshoot during the deconvolution process can sometimes occur, especially when the Single-Dish weight is too high. This non-monotonic flux convergence property can be understood as a result of the competition between the need to determine to first order all the flux (measured only with ACA), and then to locate precisely point sources (located only with ALMA).

To illustrate this problem, consider a set of 5 points sources at the corners and center of a square. Depending on the size of the square, the joint deconvolution may be fast or slow. For a square much larger than the ACA synthesized beam, the sources will be identified as point sources by ALMA, and quickly removed. For a square similar to the ACA synthesized beam, it is possible that the SNR (Signal to Noise Ratio) in the ACA data (where it looks like a point source) is larger than in the ALMA data (where each source is separated from the others). In this case, a mean flux will first be removed, and details of the structure within the square will require a larger number of CLEAN components than in the other case.

Conclusion This CLEAN-based algorithm adapted to heterogeneous arrays proved to be quite valuable and able to handle a wide variety of images. It can still be improved in several ways: e.g. by constraining the number of components to be found in each image, or by using a smoother kernel to derive the ALMA components from the ACA components, etc...

4.3.2 *uv* Plane Hybridization Technique

The method consists in combining two images, e.g. the deconvolved ALMA+SD image and the deconvolved ACA+SD image, in the *uv* plane. The FFT of the two clean images are computed, and linearly combined by selecting one image for *uv* distances shorter than e.g. 15 m, and the other for larger *uv* distances. The result is FFTed back to the image plane to produce the final image. The method has the following free parameters: the transition radius, the detailed shape of that transition, and the images to actually be combined. To avoid discontinuity, the transition shape is chosen to be reasonably smooth. The transition radius was chosen as 15 m, since beyond that ALMA data usually has better SNR than ACA data.

Another possibility is to combine clean images from ALMA+SD and ALMA & ACA+SD. By doing this, one can combine the advantages of several deconvolution methods and produce an optimum result. For example, because of the slow convergence of the joint ALMA & ACA+SD deconvolution, the deconvolved ALMA+SD image often better reproduces the long *uv* ranges, while the ALMA & ACA+SD image better reproduces the large scale structures, i.e. the short *uv* spacings. Provided the transition radius is well chosen, the hybrid image is thus better than any of these two.

In practice, the differences between both type of combinations were usually limited. The chosen combination (ALMA+SD and ACA+SD) has two significant advantages: *i*) it requires only very limited computation in addition to the deconvolution of ALMA+SD (the deconvolution of ACA+SD is much faster than ALMA+SD because of the smaller number of pixels) and *ii*) although we actually use a deconvolved ACA+SD image, the sampling of the *uv* plane up to 15 m is complete, so a linear combination of the initial visibilities could in principle be made. Noise amplification due to deconvolution should thus be very limited (see Schwarz 1978 [7]).

4.4 Relative Weight of Data

4.4.1 Single-Dish versus Interferometer

In all cases involving short spacings, the relative weight of the single dish data to interferometer data is critical. Within the restrictions imposed by the noise level, this weight is a free parameter. The usual approach is to match (as far as possible) the Single-Dish and interferometric densities of weights in the *uv* plane. This is what is required to provide a complete Gaussian beam from a truncated Gaussian weight distribution in the *uv* plane. This approach is optimal in terms of noise, since the restoring beam is usually Gaussian.

The weight distribution of the ALMA compact configuration is not yet completely decided, but is likely to have some over density of weights at the shortest spacings (15 to 25 m). Hence, when doing the single-dish weight matching, it may be better to consider the average weight density in the 25 to 50-m spacing range. For ACA, which only has 12 antennas in our study, the weight distribution is very uneven, and matching the weight density of single-dish is more problematic. The optimum result may depend on the image being restored.

4.4.2 ALMA versus ACA

Similarly, the relative weights of ALMA and ACA data can be derived from weight density in the *uv* plane. The optimal relative weights directly determine the relative integration times of each data set (ALMA, ACA and Single-Dish). In practice, a good result is obtained with ACA and the 4 12-m antennas observing 4 times longer than ALMA. Another possible approach is related to signal to noise considerations. Since there is much more flux at short spacings than on longer baselines, and since these are affected by calibration errors anyhow, it may be worth relaxing the noise level at the shortest baselines, hence diminishing the integration times with ACA. This obviously would depend on the image.

5 Understanding Fidelities

Estimating the quality of the simulated image as compared to the original model is not an easy task: the actual criteria to estimate the observation quality depends on the science that is to be done! Despite this caveat, and in order to provide a quantitative evaluation of the methods, we chose to present the results in terms of **fidelity**. Fidelity is the inverse of the relative error. A fidelity of 100 is equivalent to a 1% relative error.

Fidelities can be computed either in the image plane or directly in the *uv* plane. While the latter tell us which space scales are well reproduced, the former quantifies the visual impression coming from the comparison of the model image with the simulated one. A thorough investigation of the imaging capabilities of the instrument must take into account those two pieces of information.

5.1 Image Plane

In the image plane, the fidelity is defined as

$$\text{Fidelity}(i,j) = \frac{\text{abs}(\text{Model}(i,j))}{\text{abs}(\text{Model}(i,j) - \text{Simulated}(i,j))} \tag{3}$$

In practice, the very high fidelity points due to value coincidence between the model and simulation have to be flagged out, because they do not reflect the mean accuracy of the

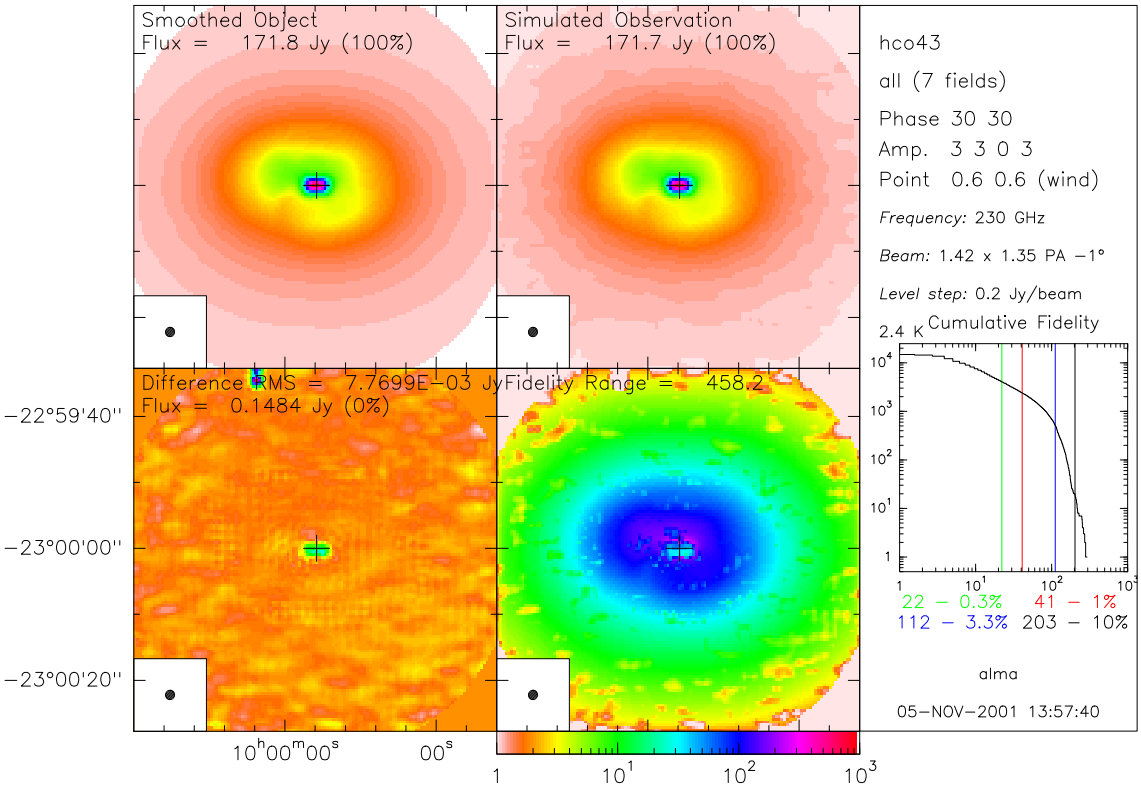

Figure 4: Image plane display of the fidelity.

simulation. For that purpose, the lowest values of the difference image are truncated, and the fidelity image is thus computed as:

$$\text{Fidelity}(i,j) = \frac{\text{abs}(\text{Model}(i,j))}{\max(\text{abs}(\text{Difference}(i,j)), 0.7 \times \text{rms}(\text{Difference}))} \tag{4}$$

The `DISPLAY` task plots the fidelity image as well as a cumulative histogram of the fidelity values, i.e. the number of pixels whose values are larger than a given fidelity. It also computes the median fidelities, taking into account only the pixels whose intensity in the initial model image is higher than 0.3, 1, 3 or 10% of the peak. Using the statistical properties of *median* values, *median* relative errors can be computed by just taking the inverse of these fidelities. Fig. 4 displays the summary of the fidelity in the image plane. It is a composite display containing:

- Top Left: the model image convolved with the CLEAN beam;
- Top Right: the simulated, deconvolved image produced by the observations;
- Bottom Left: the difference image;
- Bottom Right: the fidelity image;
- Right Insert: the cumulative histogram of the fidelity image, with median fidelity values for pixels above 4 thresholds (as fraction of the original image peak).

Those 4 median fidelities are convenient when doing a large number of simulations: they allow a global estimate of the effect of one of the parameters (e.g. pointing errors) on the simulation results. We also compute the "fidelity range" defined by L.Kogan:

$$\text{Fidelity range} = \frac{\max(\text{abs}(\text{Model}))}{\text{rms}(\text{Difference})} \tag{5}$$

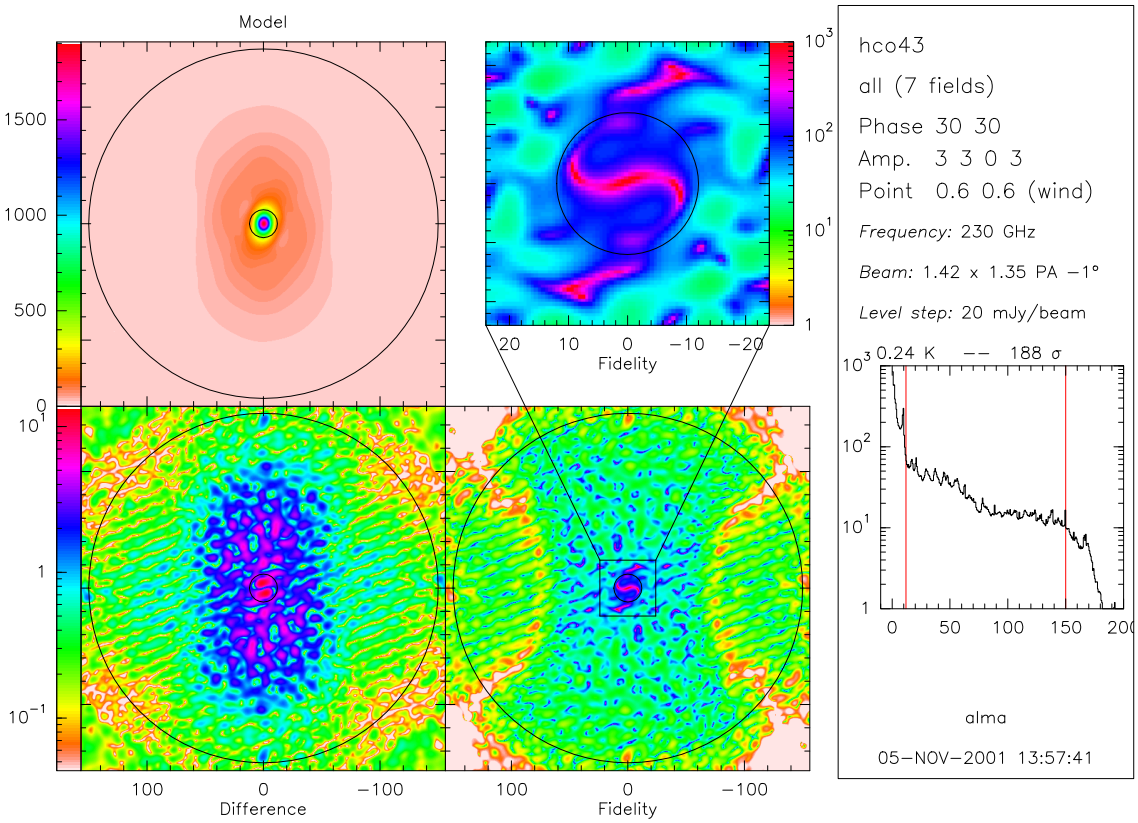

Figure 5: *uv* plane display of the fidelity.

This estimator converges toward the image dynamic range for images dominated by a single point source, and toward the median fidelity for images dominated by extended structure. In the present study, we found more convenient to use the 4 median values rather than the fidelity range to quantify the simulation results. The fidelity range is thus not used anymore in the remaining part of this memo.

5.2 *uv* Plane

The fidelity can also be analyzed in the *uv* plane. This is convenient to see which space scales are well represented, and which are most affected by errors. Typical output is shown in Fig. 5. It includes:

- Top Left: the amplitude of the FFT of the model image (convolved with the CLEAN beam);
- Bottom Left: the amplitude of the FFT of the difference between the model image and the simulated image;
- Bottom Right: the fidelity distribution in the *uv* plane, defined as the model divided by the difference; note that this is **not** the Fourier transform of the Image plane fidelity;
- Top Right: A zoom of the fidelity over the inner 20-m square of the *uv* plane;
- Insert: the azimuthal average (not median) of the fidelity as function of *uv* radius.

This type of plot is particularly useful to compare the different methods since all expected differences should only occur in the 0-15 m *uv* range. In addition, fidelity drops at longer baselines is a diagnostic of insufficient convergence of the CLEAN algorithm.

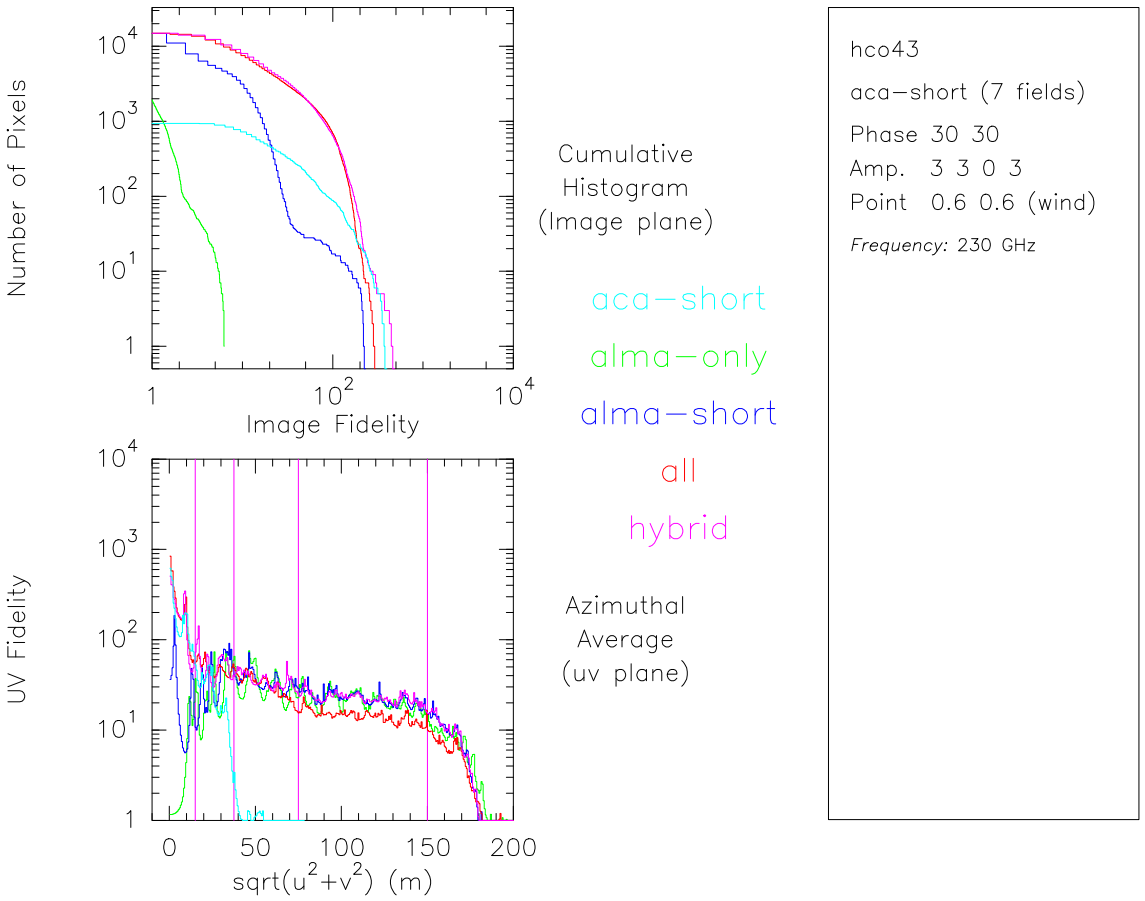

Figure 6: Overlay of the Image and *uv* plane summary displays of the fidelity for 5 different observation and deconvolution methods.

As above, fidelity medians can be computed in 4 different range of *uv* distance (0–15 m, 15–37.5 m, 37.5–75 m, 75–150 m) to allow easy estimation of trend when varying one parameter of the simulation. Some artifacts are sometimes visible in the *uv* plane fidelity, like a centered cross, and/or ring-like structures. These artifacts are related neither to the simulation nor to the deconvolution but to the presentation process: they are due to the sharp edges of the mosaic. It is possible to reduce these artifacts by applying a shallower mask. We simply chose to forget about these artifacts, since they play a minor role on the median fidelities.

5.3 Summary Plots

Comparison between the results provided by the various observation and deconvolution methods (for a same set of observing conditions) can be summarized with the two "summary" plots presented as the inserts in the image and *uv* fidelity displays. This is shown in Fig. 6.

6 Simulation Description

6.1 The Image Library

A set of test images has been used to evaluate the performances of the various observation and deconvolution methods. The images are:

- An Hα image of the M51 galaxy, with high angular resolution. The original image is 1024×1024 pixels. It contains a large number of point sources, spread over a wide area, which simulate extended emission when observed with low angular resolution. Its dynamic range exceeds 10^4. The image was provided by D.Thilker.
- An image of an HII region in M31, with complex extended structure. The image size is 128×128, and its dynamic range about 3000. The image was provided by M.Holdaway.
- A model of the HCO^+(4-3) line emission in a protostellar envelope. The image size is 256×256, and the dynamic range about 200. The image was provided by M.Hogerheijde.
- A model of a protostellar cluster, with a dozen of protostellar condensations. The image size is 512×512 pixels. The dynamic range exceeds 10^4. The image was provided by K.Tatematsu.
- A model of turbulent cloud. The image size is 512×512 pixels, and the dynamic range about 100. The power spectrum of this distribution is known. The image was provided by J.Pety.
- A model of a dusty disk around a young star, with tidal gaps created by proto-planets. The image was provided by L.Mundy.

These images have obviously different properties: spatial dynamic range from 50 to 500, intensity dynamic range from 100 to several thousands, filamentary structures in some of them, core dominated structures in others, etc... They are expected to represent a reasonable sample of the scientific problems requiring short spacing information.

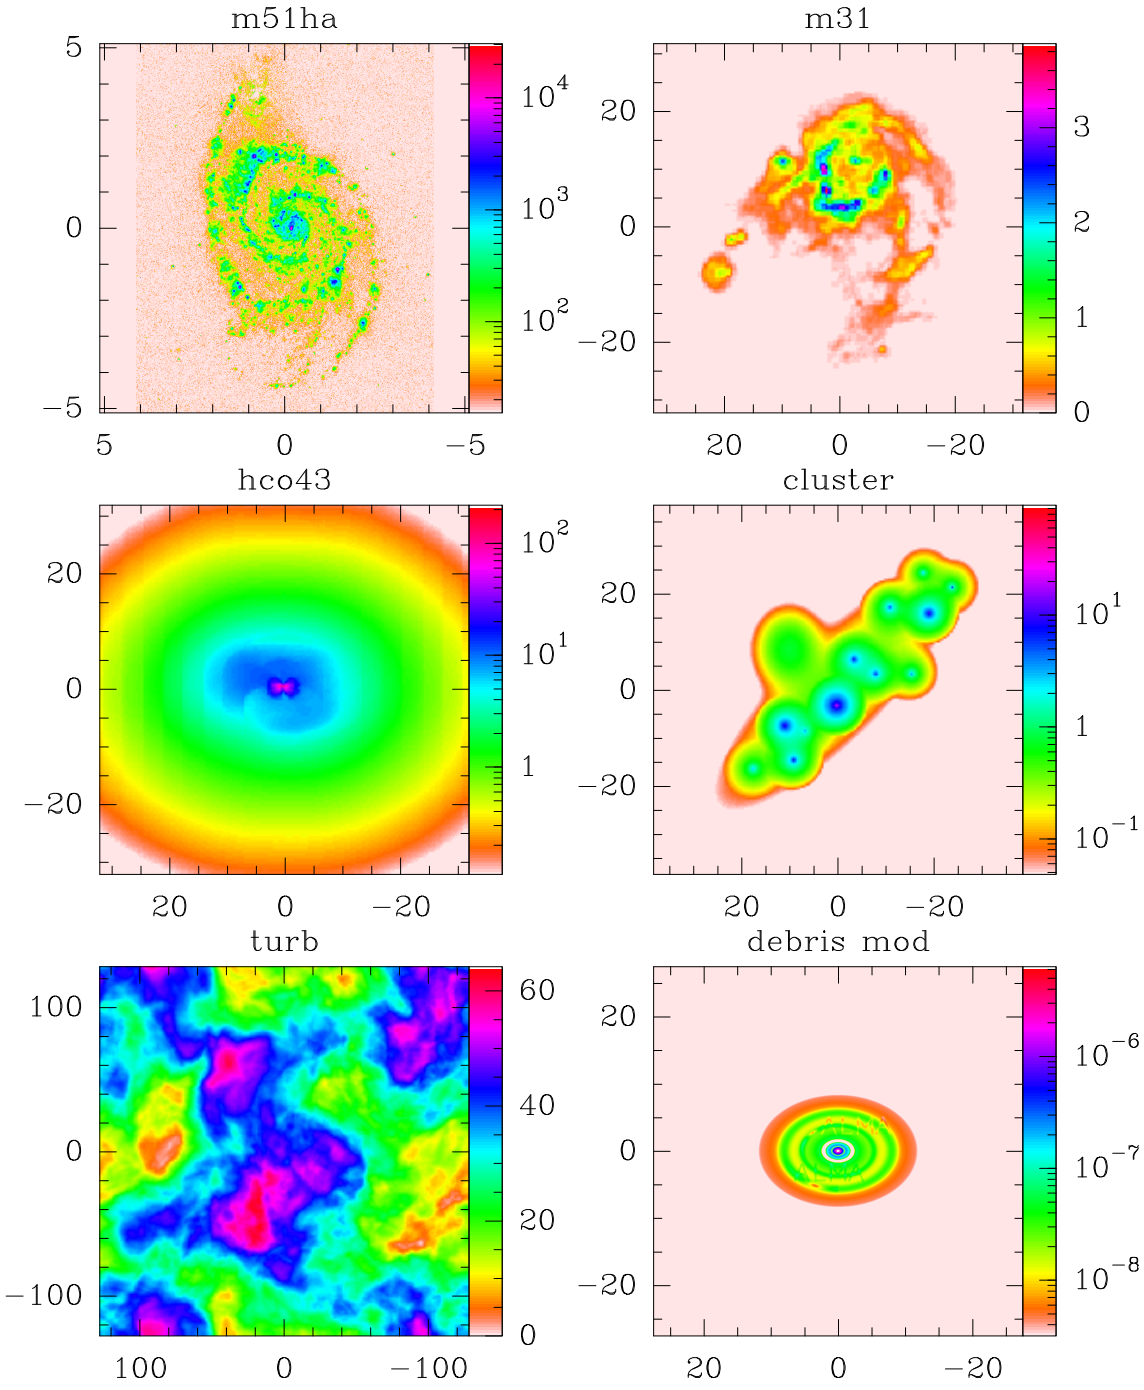

Figure 7: Images used to test the imaging properties of ALMA and ACA. Image sizes and amplitude scales are arbitrary.

6.2 Default Conditions

6.2.1 Array Layout

The array configurations used for the simulations are shown in Fig. 8. The ACA array consists of 12 7m–antennas used as an interferometer; 4 12m–antennas are simultaneously used for Single-Dish measurements. It is assumed that these 12+4 antennas are correlated together for calibration operations (Guilloteau 2001 [2]). No particular optimization process has yet been done to select the ACA configuration: it was manually selected among several randomly generated configurations. Optimization of the sidelobes level may yield better results. Note that there is a slight inconsistency between the number of ALMA antennas involved in this study (64) and the number which should actually be cross-correlated if ACA is used (60). This should however have negligible impact on the results.

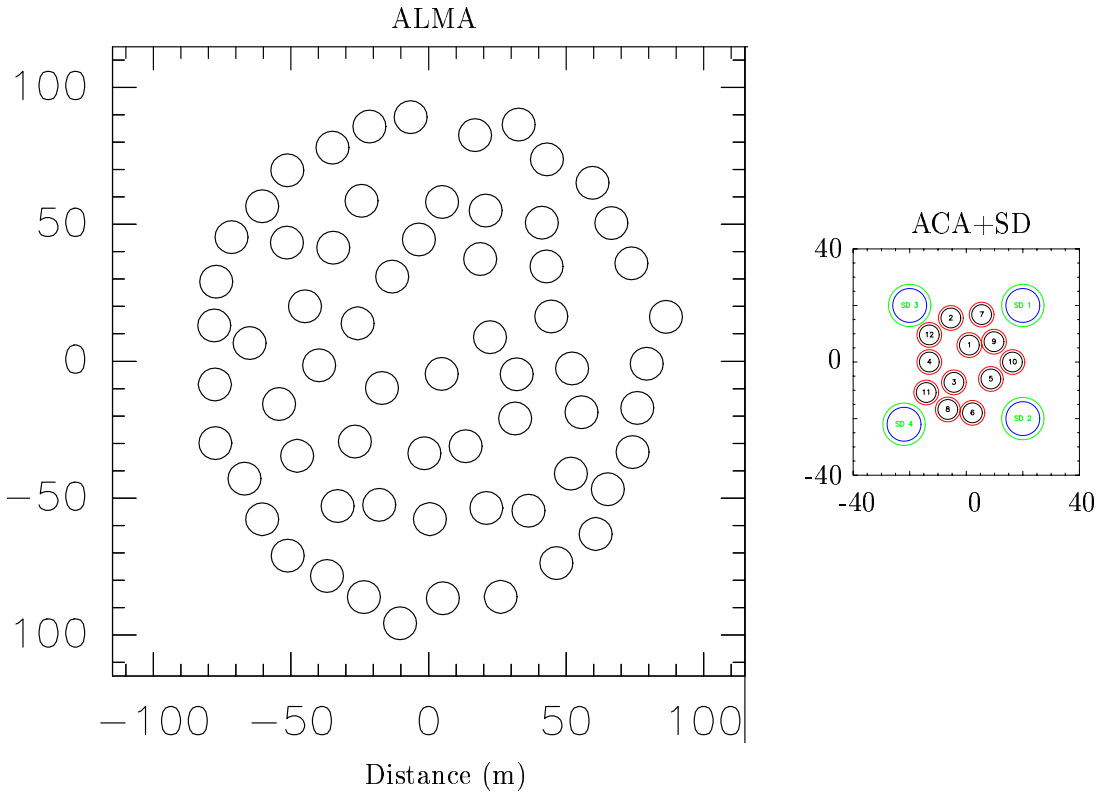

Figure 8: Selected ALMA and ACA configurations. The ALMA compact configuration was designed by M.Yun, and include 64 12-m antennas. The ACA is made of 12 7-m antennas, with close packing limited to 1.25 times the diameter, with 4 12-m antennas used as single-dish.

6.2.2 Atmosphere

We used a calibrator 2 degrees away from the source on a direction perpendicular to the wind. The phase screen was assumed to be at 1000 m above ground level, moving at 10 m s^{-1}. A "*non-connecting*" perfect WVR was assumed to be used for phase calibration, with phase referencing on the previous calibrator (Option 4 in Section 3.2.2). Note that, as shown in Fig. 3 this scheme **increases** the residual phase noise on the shortest baselines (< 30 m).

The specification of the statistical properties of the used phase screen are based on the results of the site survey. It provided a scaling exponent and a mean value on the 300-m projected baselines for the square root of the 2nd order structure function of the phase. Values to be applied to real observations depend on the observing strategy. Following Holdaway et al. (2001 [3]), we used an exponent of 0.62 on all baselines for our simulations: This may overestimate the phase noise and especially the dynamic refraction on the shortest baselines.

Using a strategy where the observing frequency is selected according to the atmospheric phase stability would result in a typical rms phase between 30 and 45 degrees at zenith. It will scale approximately as square root of the number of airmass. Other observing strategy (e.g. selecting frequency according to the atmospheric transparency) will usually result in **higher** phase noise, especially at high frequencies. However, for low frequency observations, it will still be possible to pick up observing conditions where the rms phase noise is lower than those typical numbers.

6.2.3 Observing Cycle

We simulated short observations (but not a snapshot): ALMA observes for 18 minutes, and ACA and the 4 Single-Dish 12-m antennas observe for 4 times longer (i.e. 72 minutes). The basic observing cycle was 26 seconds long: 2 seconds on calibrator, 2 seconds dead time, 20 seconds on source, 2 seconds dead time. All observations are centered around transit with sources transiting at zenith. ALMA and ACA observe the same field centers, i.e. the ACA mosaic is oversampled. For a 7 field mosaic, each field is observed 6 times with ALMA, and 23 times with ACA. Some azimuthal smoothing of the instantaneous synthesized beam thus occurs, compared to snapshot observations.

6.2.4 Range of Errors

We made computations scanning "likely" ranges of errors: 0–18 % relative pointing error (compared to the 12-m primary beam width), 0–10 % relative amplitude error, 0–60° phase errors. With 4 pointing errors, 4 amplitude errors, 4 phase errors, 6 sources, this represents 384 different simulations. We also performed simulations with one type of errors only to evaluate the effect of different error models: e.g. another set of 282 simulations were made to test 9 different sets of pointing errors. We finally made computations in "typical" high frequency observing conditions, defined as: pointing error 6 or 12% of the 12-m primary beam, amplitude error 3% (and drift 6% per hour), phase error 30°. In total, more than 700 simulations were performed.

6.3 Compared Methods

We made computations for 4 different methods:

1. **alma-only** which includes only ALMA interferometric data;
2. **alma-short** where single-dish flux has been added to the ALMA interferometric data (mosaic deconvolution of ALMA+SD);
3. **all** which results from the joint deconvolution of the ALMA and ACA+SD dirty mosaics (see Section 4.3.1);
4. **hybrid** which results from the *uv* plane hybridization of the ALMA+SD and ACA+SD clean mosaics (see Section 4.3.2).

	Intensity threshold (in % of peak intensity)															
	0.3%				1%				3%				10%			
m51ha	1	4	23	28	1	23	72	90	2	32	104	128	4	56	162	189
HCO43	1	3	83	46	1	5	132	66	1	14	260	112	1	14	506	148
turb	1	3	32	20	1	3	32	20	1	3	33	20	1	3	33	20
m31	1	139	92	119	2	160	112	145	2	186	141	175	2	223	175	211
cluster	1	103	141	145	2	141	191	179	2	178	249	213	4	268	377	286
debris	3	185	153	183	4	196	200	196	4	186	220	191	22	195	161	193

Table 1: Comparison of the median fidelities measured in the image plane for different observing modes and/or deconvolution algorithms: observation with ALMA only (**alma-only**, first columns); observations with ALMA and the 4 Single Dish antennas (**alma-short**, second columns); observations with ALMA, ACA and the 4 Single Dish and with a joint deconvolution (**all**, third columns) or with a separate deconvolution (**hybrid**, fourth columns). The median fidelities are measured for 4 different intensity thresholds. Observations simulated here are **ideal** (i.e. without noise or errors). Note that for the first 3 images, using the ACA data considerably improves the median fidelities.

	uv range															
	0–15m				15–37.5m				37.5–75m				75–150m			
m51ha	2	39	376	309	138	384	463	546	192	578	335	820	118	392	134	505
HCO43	2	15	468	545	56	104	307	145	67	151	244	213	30	59	71	73
turb	1	2	29	38	4	6	18	16	6	10	9	20	4	5	4	5
m31	2	827	496	501	23	1396	328	1265	50	1683	307	1696	29	752	103	803
cluster	2	491	484	523	94	1788	736	1833	118	1558	385	1455	65	898	153	763
debris	4	326	1545	1568	249	2315	1273	3011	663	3459	1751	5325	505	2164	995	2910

Table 2: Idem as Table 1 except that fidelities are here directly measured in the Fourier plane between 4 uv ranges. Note that for the first 3 images the use of ACA data improves the 0-15 m median fidelities by at least one order of magnitude (in the present **ideal** case).

7 Simulation Results

7.1 Error Free Case

The error free case is a useful limit *i)* to figure out the properties and precision of the simulation and deconvolution methods, and *ii)* to provide a reference. A summary is given in Tables 1–2.

A striking result is that some images are better deconvolved with ACA, others without. However, in all cases, using ACA produces a good result, while some catastrophic failures occur without it. An extreme example is the `turb` image, which is only well reconstructed with the **all** method; this image has very little power in small scale structures, and CLEAN has real difficulties in handling it. A second result is that all the differences indeed come from the 0-15 m range, which usually dominates the image plane fidelity. Extremely high fidelities can be obtained, showing that the approximations made in the simulation process are small enough for our purpose.

7.2 Pointing Errors

Simulations were made to assess the effects of pointing errors. Three different models of pointing errors were used: wind induced pointing errors (`Wind` model), thermally induced pointing errors (`Thermal` model), or both together (`Both` model). Sample results are given in Fig. 9–10 for two sources, `HCO43` and `cluster`. These results do not include anomalous refraction, which has negligible effect (see Section 7.4).

The results properties are best understood by looking at the uv plane fidelities. **all** always give better fidelity in the 0–15 m uv range, but sometimes lower fidelities than either **alma-short** or even **alma-only** at the longest baselines. When pointing errors approach 10% of the beamwidth, the three methods give essentially similar fidelities in the uv plane

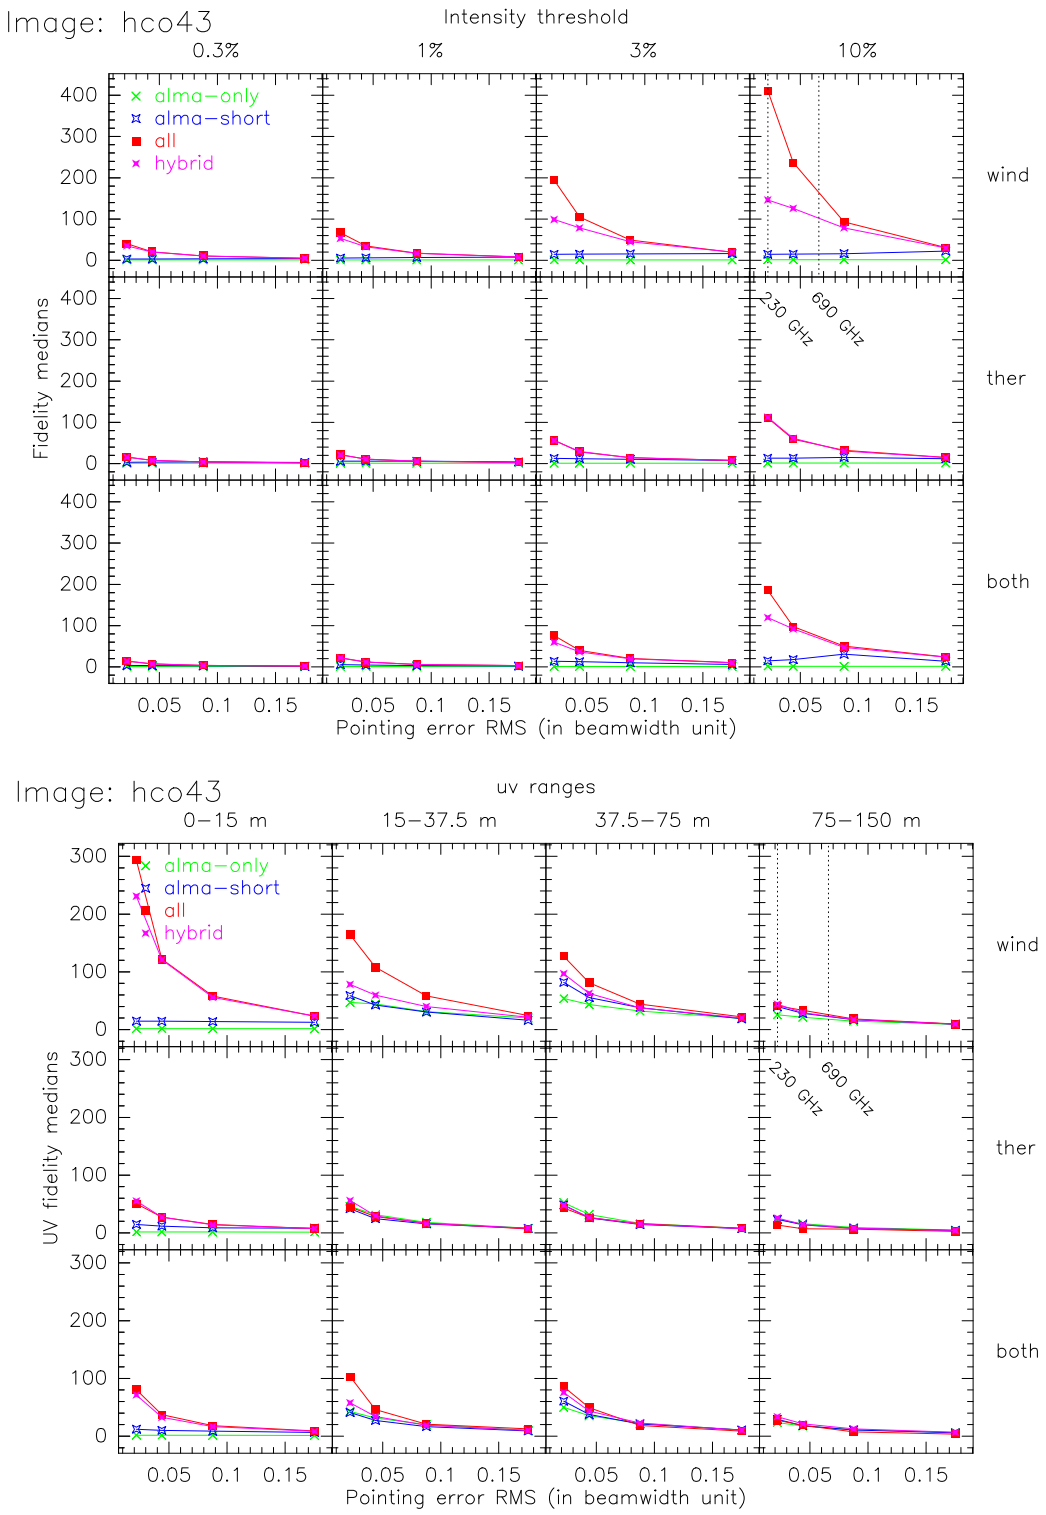

Figure 9: Median fidelities computed in the image plane (top) and the *uv* plane (bottom) as function of pointing errors. The model image is here `HCO43`.

ranges above 15 m. **all** continues to give better results in the range 0–15 m. The **hybrid** method, which was used with a transition radius of 15 m, basically combines the best of both worlds. It is only marginally worse than **all** for `HCO43` because the 15–30 m *uv* range is still poorly represented by **alma-short** in this image. In the image plane, **alma-short**, **all**

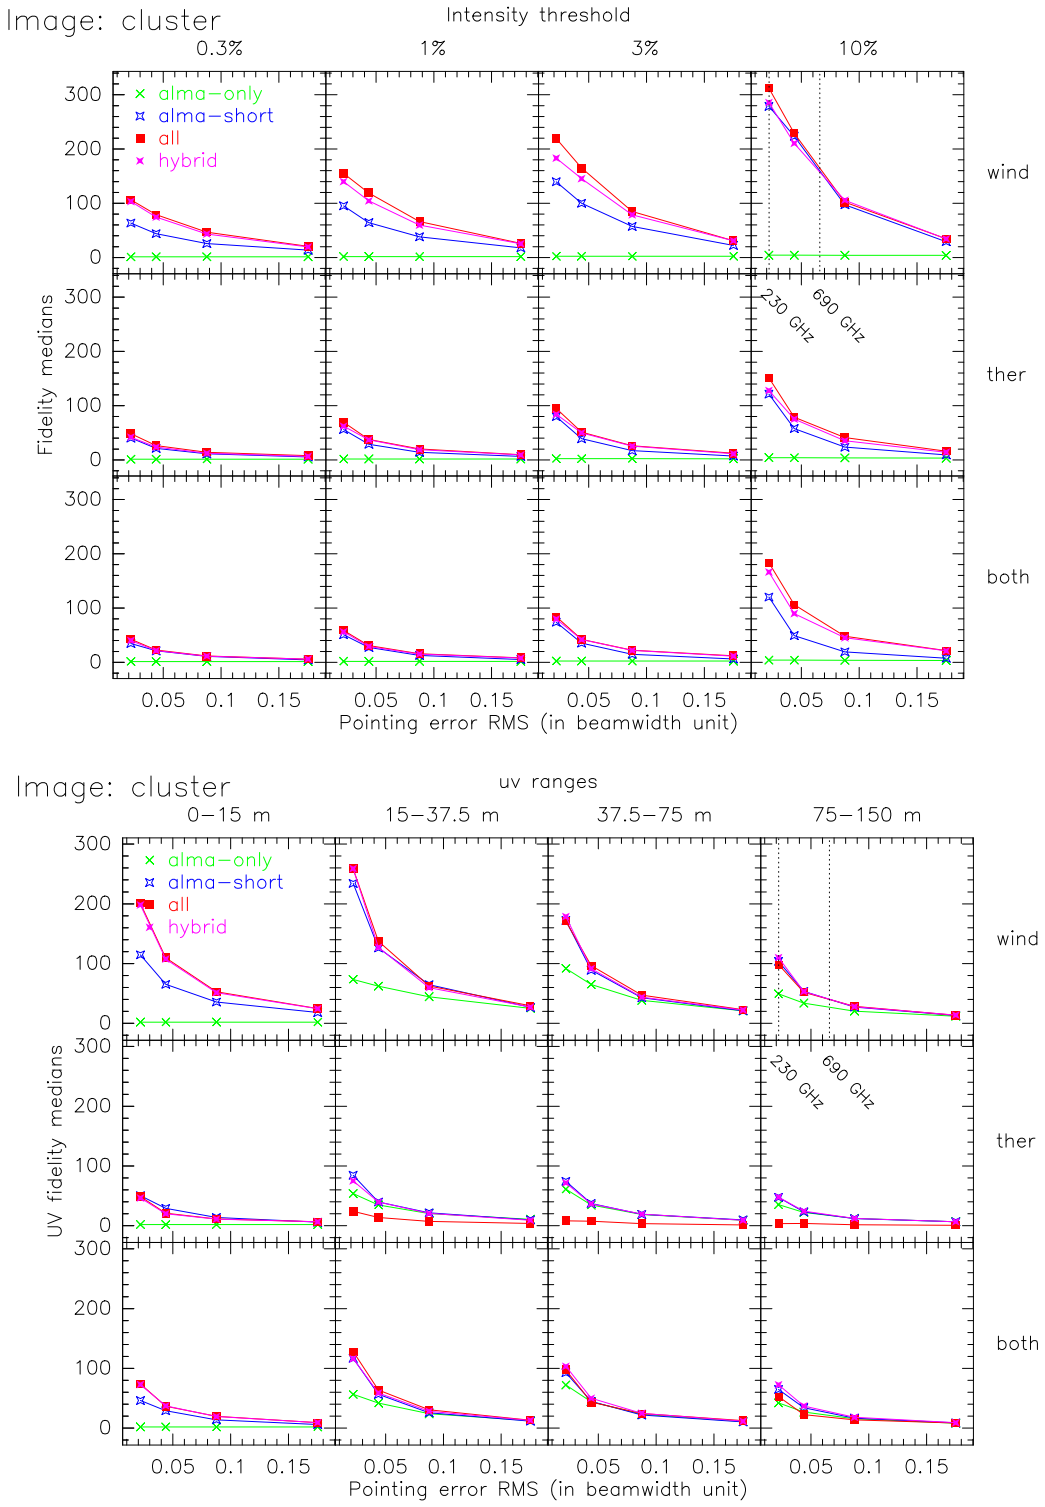

Figure 10: Same as Fig. 9 except for the model image, which is here `cluster`.

and **hybrid** converge to very similar results for large pointing errors, but **all** and **hybrid** remain superior in practically all cases.

The effect of pointing errors should be understood in a statistical sense. Because the observing duration is short, the results for any particular simulation may depend on the exact values of the generated pointing errors. In some circumstances, a pointing error may actually lead to increased image fidelity for part of an image where the deconvolution

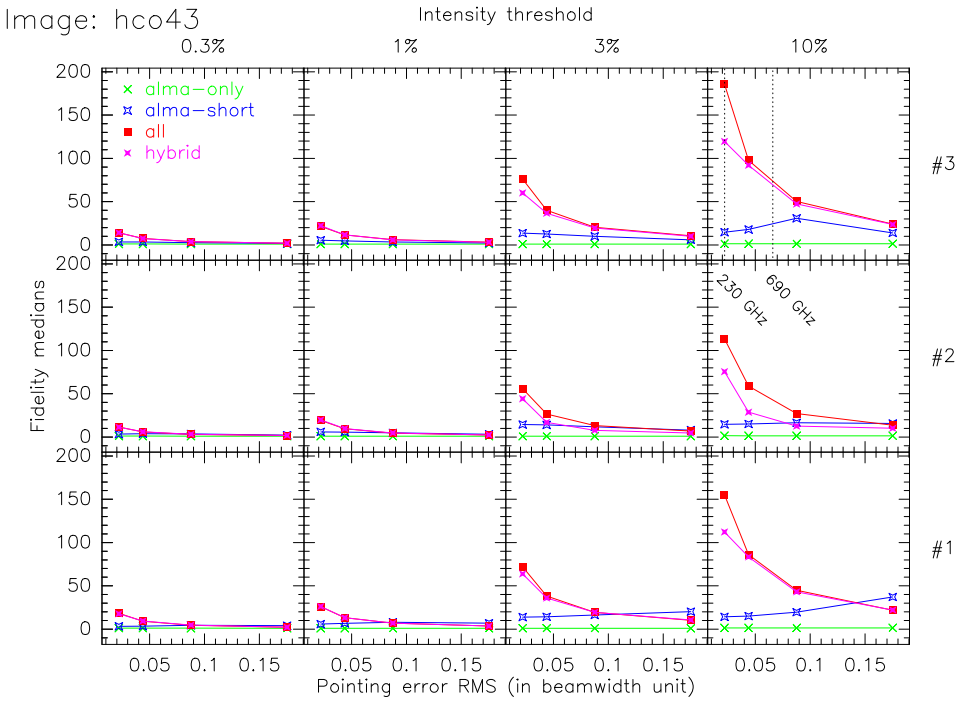

Figure 11: Image plane median fidelities for HCO43 for 3 different samples of pointing errors.

method fails in the absence of error. The median fidelity for the brightest pixels of HCO43 is particularly likely to suffer from such coincidences, because this image has only one bright peak and thus few pixels above 10% of the peak intensity. Effects of different pointing samples with the same rms value are shown in Fig. 11. It is seen that fidelities from **alma-short** occasionally improves for some particular pointing error, while the behavior of either **all** or **hybrid** is quite smooth and monotonic.

7.3 Amplitude Errors

Fig. 12 and 13 shows image plane and *uv* plane fidelities as function of amplitude errors. Amplitude errors have only been applied to Single-Dish data in Fig. 12, and to both Single-Dish and interferometric data in Fig. 13.

Even small amplitude errors on the Single-Dish data (2 %) can have dramatic effects on the fidelities in the 0–15 m range but, as expected, little or no impact at all on the fidelities at longer baselines (see Fig. 12). When amplitude errors are also applied to the interferometer data, the fidelity (in the image or the *uv* plane) drops smoothly with increasing error (see Fig. 13). Similarly to pointing errors, a single example is not sufficient to evaluate the effect of amplitude errors, and the image plane fidelities can get spuriously high in some cases.

Although the adopted observing strategy for the Single-Dish favors the **alma-short** method in presence of amplitude errors (see Section 3.3), we here find that **all** and **hybrid** remain superior to **alma-short** in all circumstances, especially in the 0–15 m range. Also note that the effect of amplitude errors on interferometric data is relatively limited if they remain below 3 %.

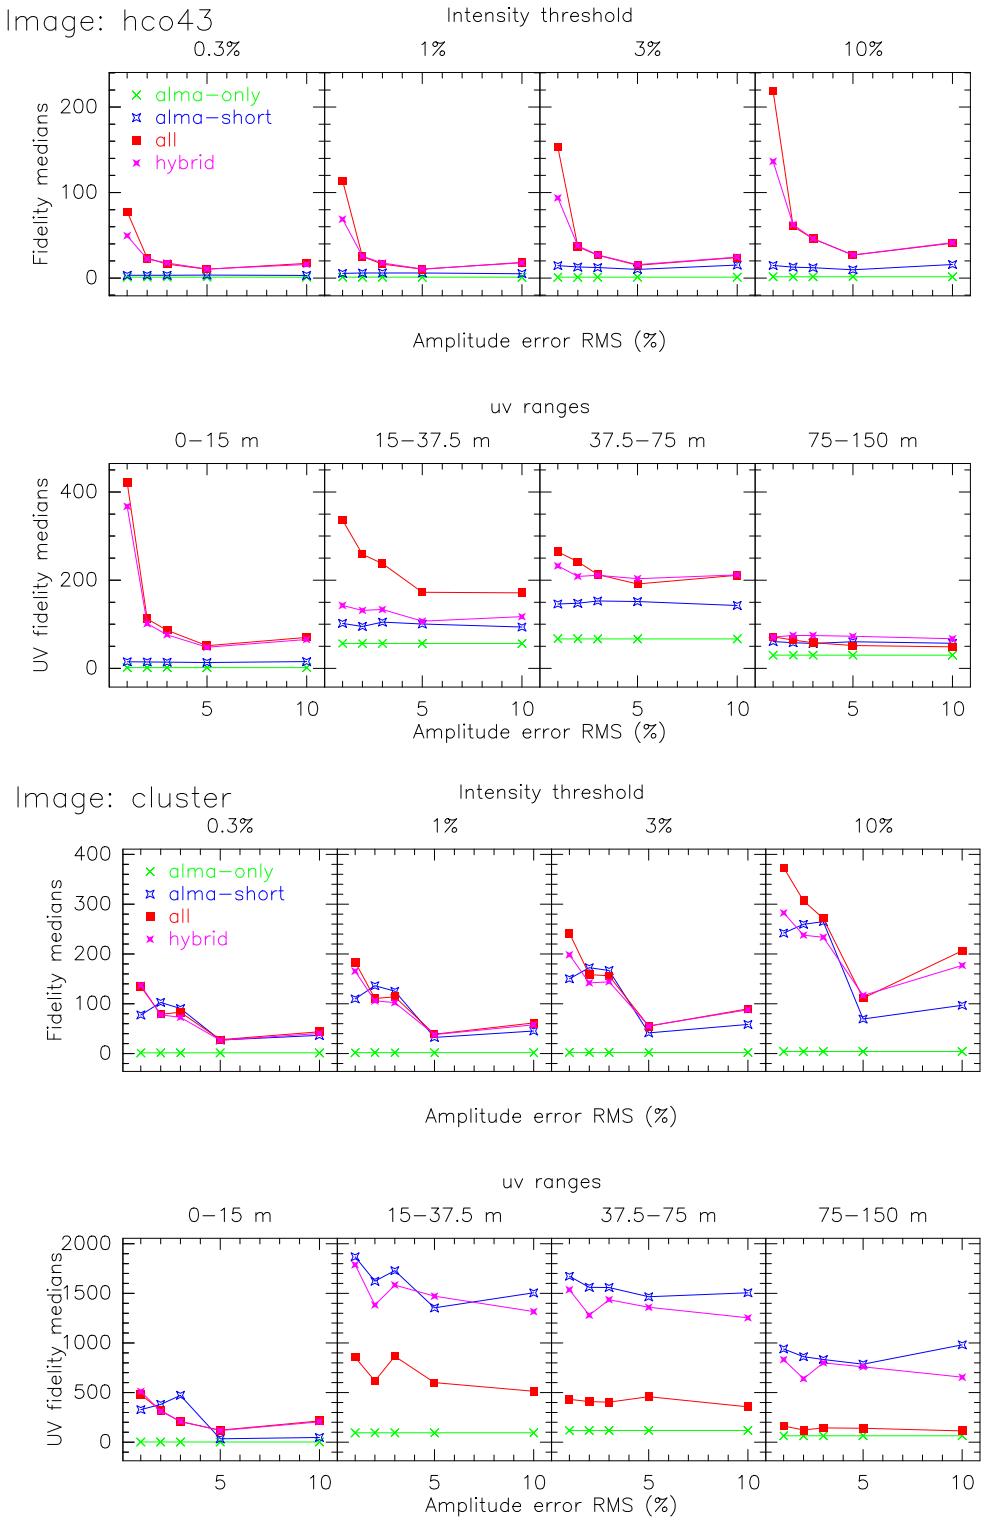

Figure 12: From top to bottom: Image and *uv* plane median fidelities as function of amplitude errors for `HCO43`, and then for `cluster`. Amplitude errors have only been applied to Single-Dish data (i.e. interferometric data is error free).

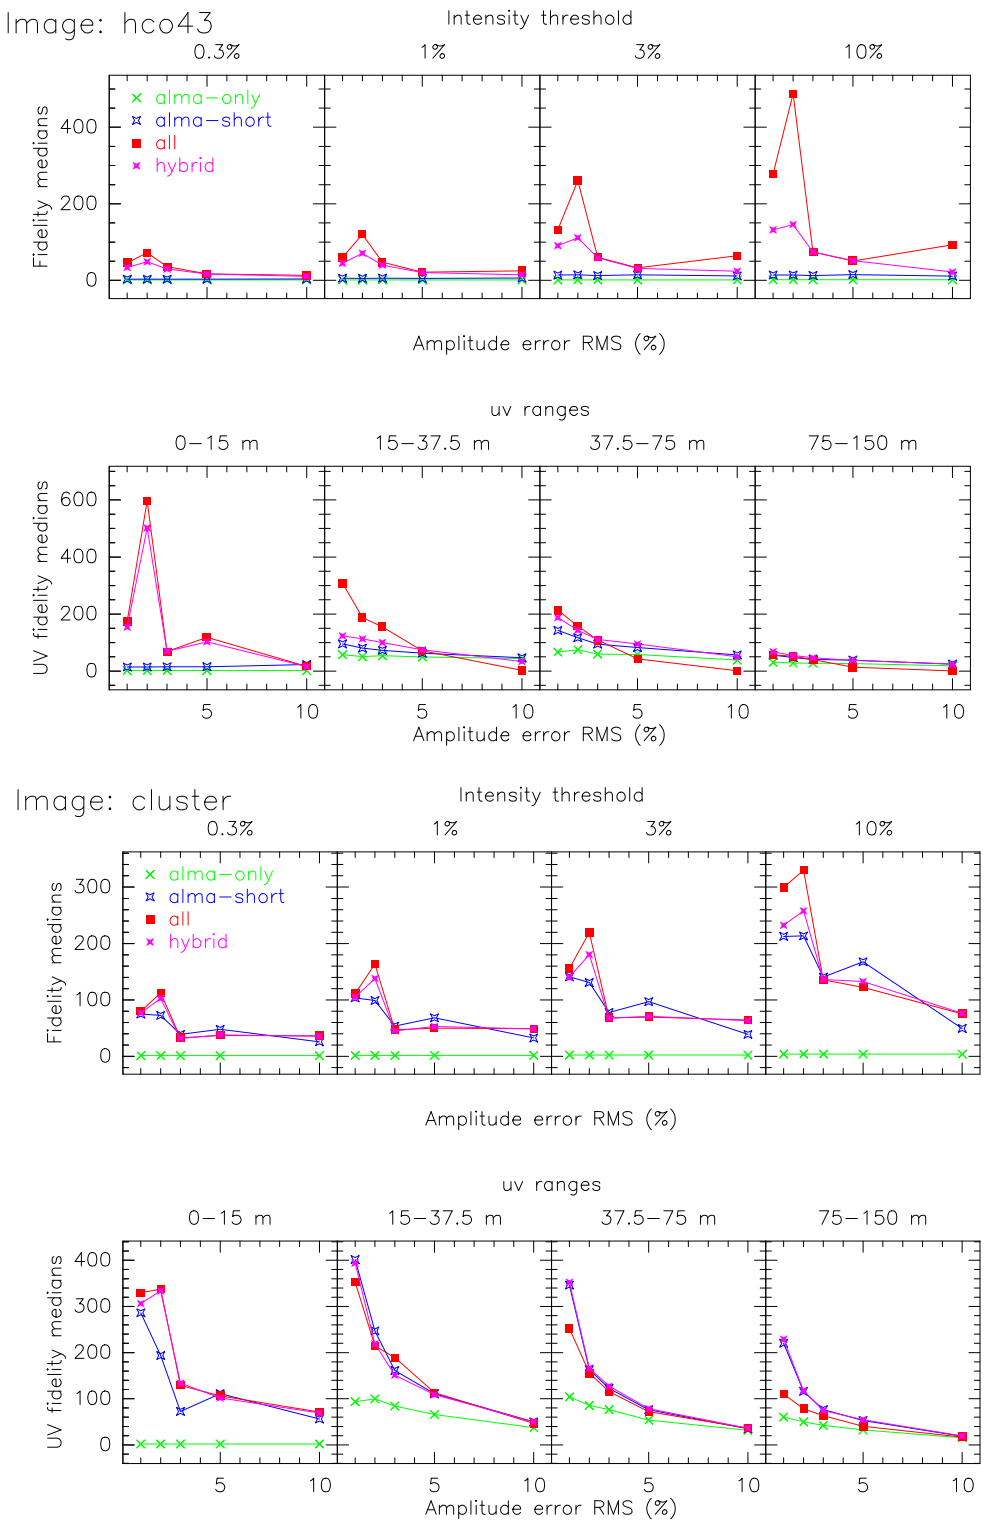

Figure 13: Same as Fig. 12 except that amplitude errors have been applied to both Single-Dish and interferometric data.

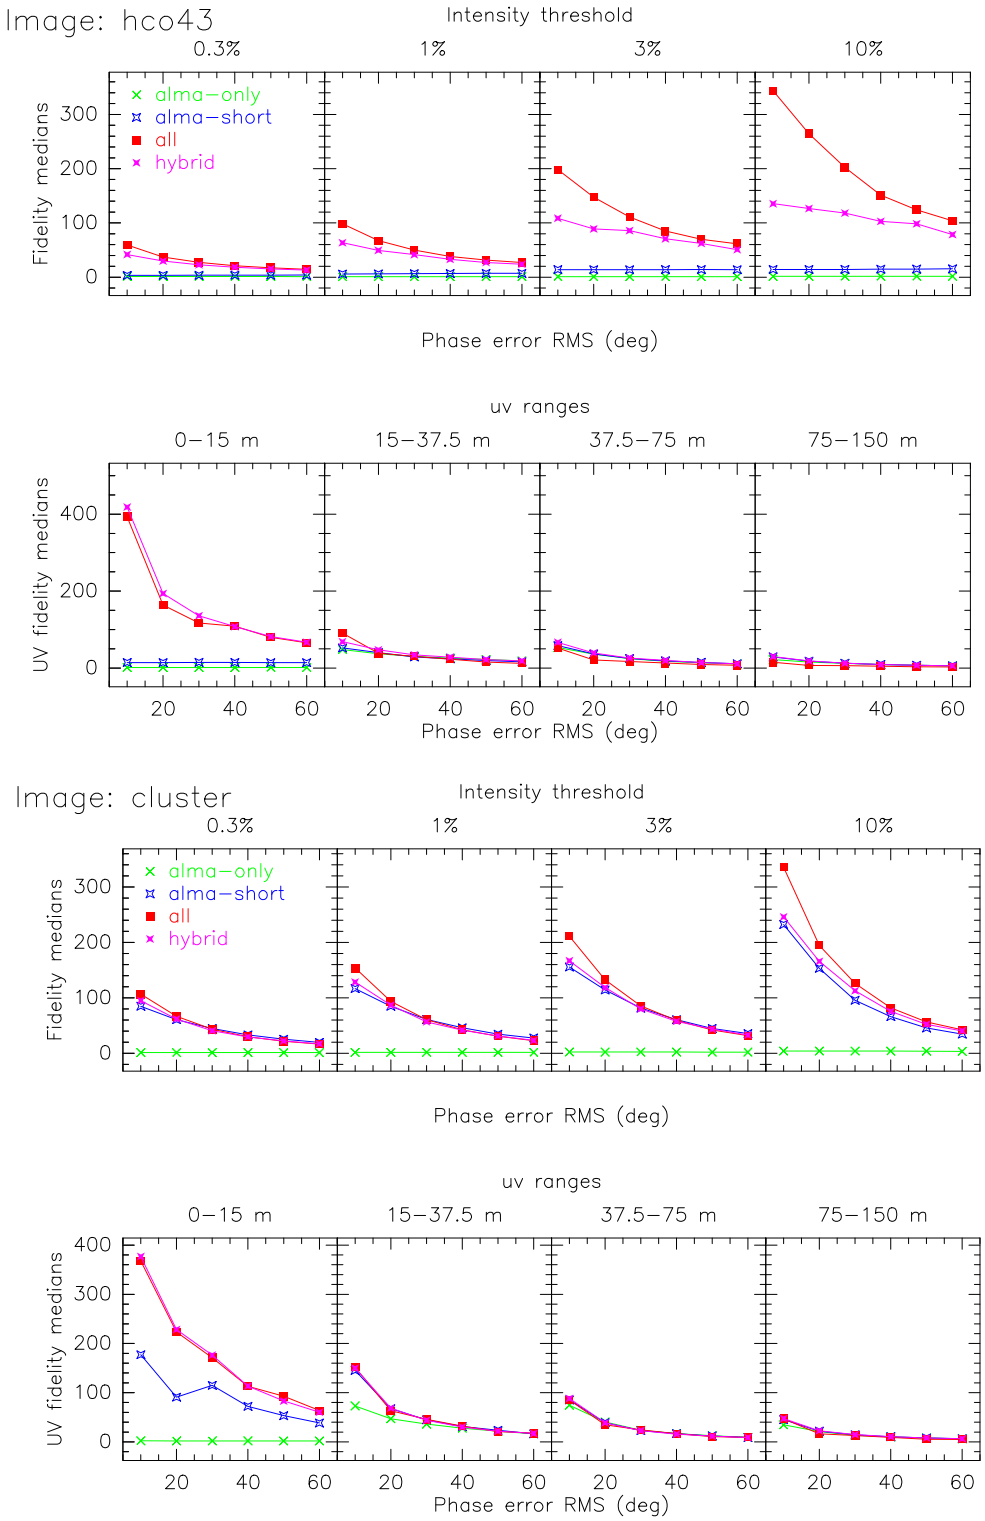

Figure 14: From top to bottom: Image and *uv* plane median fidelities as function of phase errors for `HCO43`, and then for `cluster`.

7.4 Phase Errors

Phase errors only affect the interferometer data.

Fig. 14 shows image plane and *uv* plane fidelities as function of phase errors (expressed as the value of the rms phase noise for 300 m–baselines) for two images. There is smooth

	Intensity threshold (in % of peak intensity)															
	0.3%				1%				3%				10%			
m51ha	1	5	6	6	1	20	21	21	2	30	39	41	3	60	77	83
HC043	1	4	22	19	1	7	41	34	1	15	112	75	1	15	203	111
turb	1	2	29	16	1	2	29	16	1	2	29	16	1	2	29	16
m31	1	39	33	36	1	47	40	46	2	60	52	58	2	75	66	72
cluster	1	40	47	43	2	56	69	60	2	77	96	81	4	96	126	101
debris	3	29	51	55	3	32	69	77	4	35	94	95	14	56	56	70

Table 3: Same as Table 1 except that errors are set at their typical values.

	uv range															
	0–15m				15–37.5m				37.5–75m				75–150m			
m51ha	2	39	56	48	35	41	44	43	30	30	29	30	21	21	21	21
HC043	2	15	175	183	21	41	46	52	29	33	33	33	20	22	15	23
turb	1	2	39	17	4	5	14	14	6	9	8	19	4	5	4	6
m31	2	68	85	103	30	52	41	49	28	28	29	28	21	21	16	21
cluster	2	85	142	167	54	67	66	65	36	35	37	35	23	24	23	24
debris	3	36	221	191	54	60	45	62	37	37	46	37	26	26	27	26

Table 4: Same as Table 2 except that errors are set at their typical values.

(almost monotonic) degradation of the fidelities with increasing phase error. For most images (like `cluster`), all methods converge toward the same (rather poor) result at the largest error values. **all** and **hybrid** remain however superior in all circumstances, especially in the 0–15 m range.

Contrary to pointing and amplitude errors, the average effect of phase errors is well represented by a single simulation because there are many random phase errors generated in a single observation (one random number per antenna per integration time, or about 4000 numbers for one observation).

We also checked the effect of dynamic (anomalous) refraction associated with the simulated phase screen. We found that the effect of the associated pointing error was small compared to the direct effect of the phase. Moreover, it is even possible that the scaling exponent value (0.62) used for the phase structure function actually overestimates phase errors on short baselines, where one would expect 3-D turbulence to dominate and lead to a steeper exponent (0.83). Accordingly, anomalous refraction needs only to be considered through its effect on pointing calibration on one hand, and on single-dish observations on the other hand. The impact of anomalous refraction on single-dish data will depend on the array layout (through the relative location of the Single-Dish antennas), and on the observing strategy.

7.5 Combined Errors in Typical Conditions

In real life, all types of errors are present. From the ALMA specification, we may expect a typical amplitude calibration error of no more than 3%. For pointing errors, the specification is 0.6″ which translates into 1 to 10% of the beam width for frequencies between 100 and 900 GHz. However, anomalous refraction, whose value typically ranges from 0.5 to 2″ may add up as a measurement error to this value. Atmospheric phase noise has a typical value of 30–45° on a 300-m baseline, if observing frequency is matched to the phase conditions.

We thus made a set of simulations assuming the following "typical" errors: 3% amplitude

errors, 3% beamwidth fractional pointing error (0.6″ in each direction at 230 GHz), and 30° phase noise on 300-m baseline, with a 0.62 scaling exponent. The pointing error was assumed to follow our "Wind" error model (indeed high frequency mosaic observations could be done at night where thermal effects are negligible). The results are presented for all images in Tables 3–4. An interesting property of these results is that, despite the large differences between the images in the absence of errors, image plane fidelities become quite similar with these typical errors. Image plane fidelities of 30 to 50 are found for the 0.3 to 3% intensity thresholds. Amplitude errors are negligible in this result, while pointing and phase errors have similar importance.

Note also that in the presence of these errors, the **hybrid** method basically gives the best results on all images, even if **all** was significantly better in the absence of errors. The **hybrid** method also gains an advantage over **alma-short** for all images.

7.6 Other Tests

7.6.1 Primary Beam Errors

Our model of the primary beam assumes a truncated Gaussian in the *uv* plane (i.e. a Gaussian illumination distribution up to the antenna radius) and a perfect Gaussian (with matched beamwidth) in the image plane. While this is inconsistent from a theoretical point of view, this corresponds reasonably well to what is occurring in real-life, where a more detailed model of the primary beam might be difficult to implement.

We also performed (though accidentally...) a series of computation where the primary beam was truncated to its 10 % level in the image plane. While this may be a rather large error on the knowledge of the primary beam, these computations allowed us to find out that the **all** or **hybrid** methods are much less sensitive to errors on the primary beam than **alma-short**.

This should be kept in mind when considering high frequency mosaics, where the wavelength is closer to the precision of the surface accuracy (i.e. the primary beam will only approximately be known because of surface errors and varying shape of the antenna dish). A systematic (although simple) investigation of errors on the primary beam could be made with our package, which offers a scaling parameter for the beamsize.

7.6.2 Mosaic Oversampling

We tried oversampling the ALMA mosaics by factors 1.49 and 1.73. The corresponding mosaics had 13 and 23 fields respectively, and covered a somewhat larger region than the original 7 fields mosaic. We did not find any significant improvement in the **alma-short** results when doing so. There was some improvement in the **all** results for long baselines, presumably because oversampling effectively gives measured visibilities at more *uv* points. It however did not improve at all the recovery of short spacings.

7.6.3 Other Combined Deconvolutions

We tried a joint deconvolution of ALMA+SD with ACA+SD, rather than ALMA-Only with ACA+SD as in the **all** method. This produced significantly worse results.

On the other hand, combining the results of **alma-short** with those of **all** to obtain the **hybrid** results improved the image fidelity on the most difficult cases (the `HCO43` image). However, it was also necessary to increase the transition radius to obtain this improvement. We believe that our original hybrid (ALMA+SD and ACA+SD) provides the advantage of much faster computation, and a robust default value for the (a priory adjustable) transition radius.

8 Comparison with Other Work

The computations presented here represent the most extensive set of imaging simulations performed for ALMA. Only part of the results have been presented in this document, but the complete set can be accessed from the ALMA Web site at IRAM[3].

We only used CLEAN-based deconvolution methods. However, M.Holdaway (NRAO) and K-I.Morita (NAOJ) performed similar, though less extensive, simulations using MEM-based deconvolutions (Maximum Entropy Method). They also use a different method to derive visibilities. Their simulation process is more accurate in the absence of errors, because of the handling of the pointing problem. But both simulations give very similar results in the presence of errors. In particular, it appears that the difference in behavior between the various images is not related to the use of a particular deconvolution tool, but rather is an intrinsic property of the image.

K-I.Morita also performed simulations including thermal noise, and found that the behavior depends in a rather complex way on the noise level. Basically, MEM requires a SNR of at least 5 on ACA visibilities to perform properly, and can catastrophically fail otherwise. We have not yet performed simulations with a varying noise level. However, in our **hybrid** process, the operations performed to recover the short spacings are essentially linear (because the *uv* coverage is complete, see Schwarz 1978 [7]). We thus expect this method to be stable with respect to noise. On the other hand, it is possible that the **all** method be less stable in this respect.

M.Holdaway suggested that adding a "guard band" of fields around the useful mosaiced region significantly improves MEM imaging in the **alma-short** case. This may be understood when the emission falls to zero at the image edge, since it would provide a reference value for MEM. Whether this result also applies to CLEAN remains to be studied. Moreover, it is unclear that this improvement also occurs in a complex area where emission extends well beyond the mosaic edge.

9 Conclusions

In summary, we found that:

1. The homogeneous array, even with the inclusion of zero spacing (ALMA+SD), fails to produce reliable results on some type of images. This failure is not so much related to unavoidable observational errors (phase noise, pointing or amplitude errors), but rather to the amount of smooth extended structure present in the image.

2. The addition of data from a compact array of 12 7m–antennas (ACA) allows one to obtain reliable results in all circumstances. Although the newly developed joint deconvolution (ALMA & ACA+SD) sometimes produces the most accurate image, in presence of typical observational errors, the simpler and much faster *uv* hybridization method (ALMA+SD with ACA+SD) gave similar or even better results on all test images.

3. The inclusion of ACA data allows one to produce images which are much less sensitive to primary beam errors than ALMA+SD. Although we did not quantify this advantage, it should not be forgotten since primary beams may actually not be well known, especially at the highest frequencies.

[3] http://iram.fr/~alma

4. We found that the effect of pointing errors due to anomalous refraction is much less significant than the direct effect of phase errors on the visibilities when standard (i.e. observed) parameters of the atmospheric model are used.

5. The data processing remains simple, and does not require any significant additional computing power.

Although the above results are based on CLEAN-based deconvolution techniques, it is important to stress that independent work based on MEM deconvolution have confirmed the main conclusions.

We thus conclude that the addition of the ACA brings a significant advantage to ALMA in typical observing conditions, and provides the following improvements to the ALMA project:

1. A more robust imaging process.

2. Results more immune to pointing and primary beam errors.

3. A practical advantage for the short spacing measurements: since only 4 12-m antennas are used, it may be easier to equip these antennas with optimized hardware for single-dish. This can include wobbling subreflectors, receiver optimized for stability, specific hardware for anomalous refraction correction such as a scanning water vapor radiometer, etc... It may also be possible to select the best 4 antennas among the 64 available ones.

In addition to these basic improvements, the ACA provides an array available in a compact configuration at any time for high frequency work, whereas the 12-m antennas of ALMA may be spread in long baseline configurations when observing conditions suitable to high frequency work occur. Indeed, at 900 GHz, ACA provides an angular resolution of $1.5''$, quite comparable to that provided by the compact configuration of ALMA at 230 GHz.

Additional remark ACA certainly results in additional complexity in the operations, since it basically represents another array with its own observing program. It also induces some more complexity for construction and maintenance, by introducing a different type of antenna and receivers. Using the same antenna mount and the same focal ratio will limit, but not suppress the differences.

Future work Although fairly realistic, our simulations are amenable to several extensions and improvements:

1. The effects of thermal noise should be checked.

2. The indirect effect of dynamic refraction on pointing error measurements could be accounted for more accurately. The properties of the phase structure function imply some correlation between the pointing errors introduced by dynamic refraction on the various antennas. The amount of correlation depend on the power spectrum of the phase, and also on the array layout.

3. The effects of primary beam errors could be incorporated and studied.

4. Single-Dish observations could be simulated in a more realistic way. This would likely improve results based on the ACA+SD mosaic (the **hybrid** and **all** method), and degrade results of the **alma-short** method.

Because of the robustness of the results obtained with ACA, we do not believe any of these additional tests will result in a significant change of our conclusions.

Finally, the simulator developed in the course of this project may play a role in the future of ALMA by helping astronomers to evaluate ALMA performance on their data. Its speed (about 10 minutes per image on today's computers) makes it amenable to extensive use, and it can readily be converted into a fully automated pipeline, taking a model FITS image as input, producing a "likely" output image and estimating expected image fidelity. Such a tool could be important to help astronomers select the best observing strategy.

Acknowledgements

We thank M.Holdaway for his critical screening of our simulation results, all ACA "tiger" team members for their comments, F.Viallefond for his nice ideas about the atmospheric phase screens and H.Wiesemeyer and V.Pietu for efficient help during some parts of this study.

References

[1] B. G. Clark. An efficient implementation of the algorithm 'CLEAN'. *A&A*, 89(3):377–378, 1980.

[2] S. Guilloteau. Single-Dish and short baselines. ALMA Memo XXX, IRAM, 2001. In preparation.

[3] M. A. Holdaway and J. Mangum. Relative pointing sensitivity at 30 and 90 ghz for the ALMA test interferometer. ALMA Memo 373, NRAO, 2001.

[4] K.-I. Morita. Wide field imagings with the Atacama large millimeter/submillimeter array. ALMA Memo 374, NAOJ, 2001.

[5] J. Pety, F. Gueth, and S. Guilloteau. ALMA+ACA simulation results. ALMA Memo 387, IRAM, 2001.

[6] J. Pety, F. Gueth, and S. Guilloteau. ALMA+ACA simulation tools. ALMA Memo 386, IRAM, 2001.

[7] U. J. Schwarz. Mathematical-statistical description of the iterative beam removing technique (method CLEAN). *A&A*, 65:345–356, 1978.

[8] D. G. Steer, P. E. Dewdney, and M. R. Ito. Enhancements to the deconvolution algorithm 'CLEAN'. *A&A*, 137(2):159–165, 1984.

[9] M. Yun. An imaging study for ACA. ALMA Memo 368, NRAO, 2001.