8.1 Science

1. Additional data reduction/analysis complexity (10-20%?)

We estimate 1-2 additional FTE. Explanation as follows.

For mosaicing, the processing time is dominated by the FFT's you do, which is proportional to:

- the number of iterations (ie, dynamic range required, or complexity of the source -- difficult to quantify),
- the number of pointings on the sky
- the number of different primary beams (this is actually difficult to quantify, as *we will usually be able to get away with the assumption that Vertex & AEM antennas have the same beam* -- only for super high fidelity/ dynamic range observations will we need to make the three different primary beams)

So, a complete guess -- that 80-95% of observations can assume that the beams are all the same. The remaining 5-15% of observations will take ~3 times more CPU time than if they were from the same antenna. We think this means that our imaging/computing costs for mosaicing (already only like 30% of our observations) will go up 15% to 45%. We probably can't degrade this to account for the fact that only 30% of observations are mosaiced, otherwise the pipeline will get clogged during compact configuration (ie, we design the pipeline to meet the busiest demand).

In order to firm up that number any more would require extensive simulations to pinpoint what sort of Dynamic Range/Image Complexity would require the multiple-beam imaging, and also a detailed knowledge of the distribution function of observations. I think there is significant uncertainty in the last part. My guess is the best thing to do at this point is to implement these algorithms: make realistic model voltage patterns for the two dishes and get AIPS++ to swallow this. I think AIPS++ will actually have a difficult time with it, as there is still a lot to be done in this direction.

For polarization:

Wide field polarization observations can happen in any configuration, but will be more likely in smaller configs. Suppose that pol observations will occur ~10% of the time, wide field will be less. Holdaway's wide-field pol leakage correction scheme is actually faster than the mosaicing (it is only one iteration of mosaicing), and the pol mosaicing itself will probably not require use of different beam models, SO, I would argue that the extra computing costs for the extra wide field pol work we need to do will be only a few percent.

It increases the amount of bookkeeping in the software (carrying around V1, V2 voltage patterns or $B1^2$, $B2^2$, and B1B2 beams) and there are storage and computation issues. Therefore the upshot is probably about 1-2 total FTE-years extra computing effort that you might have to spend earlier rather than later, and maybe a 25% or so hit on effective data rates (which will occur during gridding and degridding stages) which you either have to spend extra money for computing or eat in throughput on the science cases that need it (relatively high dynamic ranges for filled beam or mosaics).

The thing to do is to IMPLEMENT this stuff. Starting with the VLA would be the right thing. And then we can extend to ALMA via simulations.

2. Additional **Commissioning complexity** (initially doubles effort, later minimal impact)
The commissioning of the new antennas occurs when the construction staff is beginning its rampdown, operations its rampup. I believe it would require extension of the period of employment for one construction period science IPT person during 2008-9, alternatively another operations person or accelerated staffing.
Estimate: 2 FTE.
3. Additional Science support for AIVC required (?)
Support needed for pointing, holography, initial radiometric tests etc. on two antennas: 0.25-0.5 FTE
4. Dynamic range, polarization, fidelity? See following discussion.

DISCUSSION

At the time of the Request for Proposals of ALMA antennas (1999-Mar-30) a goal of the prototype antenna design was to produce a design "optimized for production of a quantity goal of 64 units, taking advantage of economies that may be realized by maximum duplication and standardization of parts…" (Section 3.1.2.1(c)). Two prototype antennas were contracted by the project, of substantially different design. Tests of the prototype antennas show that both meet specifications (**Joint Antenna Technical Group Test Results** April 15, 2005). Several existing arrays, the CARMA and Nobeyama Millimeter Arrays, contain (or have contained) at least element of somewhat different design from other antennas and, of course, VLBI has used heterogeneous arrays before the advent of the VLBA. The different quadruped structure on the dissimilar NMA element resulted in different, difficult to calibrate, performance, and VLBI has concerned itself only with sources much smaller than the primary beam. Here we consider some possible science implications of different designs but constant aperture for antennas within the ALMA array, which is specified to provide high dynamic range images (ALMA Bilateral Agreement, Annex B) over a field of view much larger than the primary beam, and to perform to the highest frequency limits provided by the atmosphere at its exceptional site.

Quadrupod Configuration

We have examined the issues put forth by the ASAC (see Appendix). The first of these to examine is that arising from the rotated quadruped geometry between the two designs. Holdaway has calculated these, using shadowing estimate from ALMA Antenna Group Report #40 (2002-03-28) for Vertex, a calculation from Robert Lucas for Alcatel. Considerable discussion has occurred within the Science IPT and beyond, including Darrel Emerson, Mark Holdaway, Rob Reid (Penticton) and Dick Crutcher. The preponderance of opinion is that as long as antenna specs are met, the impact of the crossed feedlegs on ALMA's science output will be in added data reduction complexity, as accounted above.

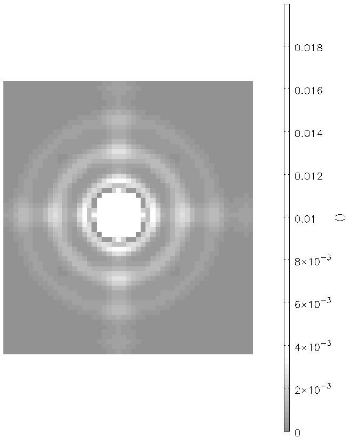

Vertex Antenna Beam

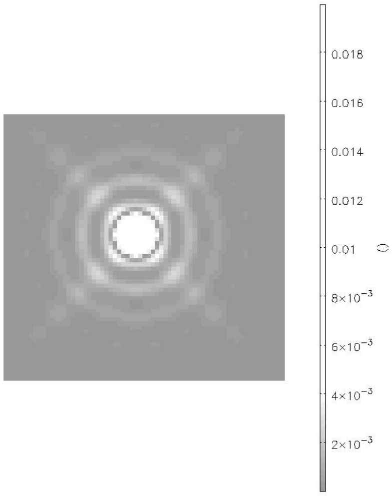

AEC Antenna Beam

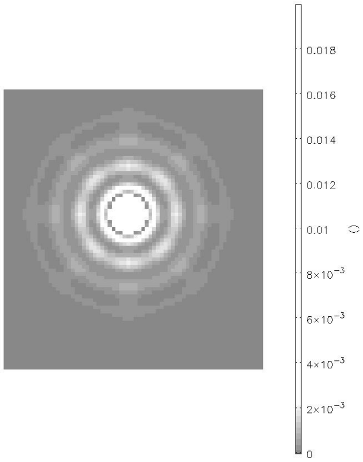

Beam on baseline with Vertex and AEC antennas

The beam profiles above were calculated with a 12 dB taper in the Primary Beam at the edge of the dish, and a perfect surface. Holdaway formed the three primary beams and looked at the RMS difference between the beams divided among primary beam regions of similar Primary Beam level. The fact the the Vertex - EIE rms is essentially twice the Vertex-Cross RMS means that these errors are systematic, but they are still very small. Still the largest error is at the 0.3% level -- a factor of 20 below the 6% specification.

The beams have been normalized to have a peak of 1.0, yet the integral of the Vertex and EIE beams differ by 0.4% -- if left uncorrected, this would translate into something like a 0.4% flux scale error. We deem this insignificant.

Common Mode Errors

For identical antennas, some effects occur in common with all antennas and common mode errors often cancel. For different antennas, they may not. Some scientific experiments may be sensitive to non-cancellation, for instance Very Long Baseline Interferometry, in which ALMA is used as a phased array. If the common mode errors do not cancel, when ALMA is phased as a single VLB element, the gain of the phased array is reduced and sensitivity in the VLB array is reduced.

Gravitational Sag

Specifications should be tightly monitored, but we believe that as long as both antennas meet those specifications, scientific performance should be adquate.

For a centro-symmetric telescope design, gravity will affect pointing through focal length changes, and through elevation terms in the pointing equation. Guilloteau has cited one example of this—the pathlength changes as the focal length of an antenna changes under the effect of gravity cancel out in the Plateau de Bure antennas, leaving a residual effect more than one order of magnitude below the direct effect. This may also occur with the ALMA antennas, build to identical specifications but differing in detail of design. These pathlength changes affect instrumental phase, and differential phase changes can cause imaging artifacts. As D'Addario and Cernicharo have noted, the error cancellation becomes increasingly important as we move to mm and sub-mm wavelengths, since the errors on the individual antennas (before cancellation) are larger in units of wavelengths or beamwidths than they are at cm wavelengths
Regarding the surface changes with ambient temperature changes, the AEG found that:

Surface RMS changes with ambient temperature from holography:

* VertexRSI: ~0.6-0.7 micron/K.
* AEC: ~0.8 micron/K.

Both deformations had a high degree of structure (like BUS segment print-through for VertexRSI, large-scale 45-degree plus inner-ring print-through for AEC).

Prototype antenna focal length change due to ambient temperature changes:

* VertexRSI:
    o 34 micron/C from holography
    o 36 micron/C from radiometry
* AEC:
    o 14 micron/C from holography
    o 20 micron/C from radiometry

In the Overview of the prototype antenna performance the testing group pointed out that:

"The specification states that a refocusing will be done every 30 minutes, whereby it is assumed that the ambient temperature will change by not more than 1.8 K/30 min. Thus a focus change, and hence a path length variation of 100 micron (twice the focus change [for the VertexRSI antenna]) between two focus calibrations can occur. This is a highly repeatable error which will be almost the same for all antennas [of the same design] in the array. By measuring the [ambient] temperature it can be quite accurately corrected, if desired."

Pointing

We conclude that differences in pointing will not be a significant problem.

. Pointing specifications for the ALMA antennas are quite tight and were set basically to provide excellent imaging at frequencies below 300 GHz for mosaicking with 8m antennas, later scaled for 12m antennas. Gravity affects pointing through elevation terms and may differ between different antenna types according to the details of the antenna structure. Differential pointing may also be calibrated between antennas, but errors will cause imaging artifacts, particularly on mosaicking, as it was the mosaicking specification which led to the pointing specification (e.g. MMA Memo No. 61, which shows that for 8m antennas operating at 230 GHz doubling the pointing error from the specification halves the dynamic range of the images. Since that specification was determined for imaging at 300 GHz and below, artifacts can be expected to be more serious at higher frequencies. In addition to the common mode pointing error, there will be baseline errors between two antennas of different types which do not cancel, but which would be minimal between antennas of the same design.

Pointing also depends upon weather, particularly with regard to the thermal characteristics of the antennas, which may vary with location over the 18km extent of the largest array. These characteristics will, of course, be minimized for a homogeneous antenna set. For a heterogeneous array, they will be minimized when the array is in its most compact configuration. For the larger configurations, mosaicking is not so important a science driver as in smaller configurations; furthermore the thermal variations may be more extreme from antenna to antenna, possibly negating some cancellation of common mode thermally induced pointing errors. Another important aspect of weather is the wind, which may affect pointing similarly for a homogeneous array in close configurations so that mosaicked images may be obtained with good dynamic range. Although winds are fairly constant in direction at Chajnantor, the extended array may place antennas in regions where wind directions differ slightly owing to local effects; this will produce non-cancelling pointing errors which in turn affect imaging, particularly for more extended sources. Differing resonance frequencies for antennas with different designs could result in different wind responses, lessening cancellation of errors even in compact arrays. Both thermal and wind effects will result in deteriorating image quality at higher frequencies, particularly for an array of a heterogeneous antenna design.

Fiber Length

Further thermal effects will affect imaging through changes in fiber length. Dick Thompson notes that "the effective length of the fiber is dominated by the run up the antenna (see ALMA Memo 443), and for this the length and degree of thermal shielding may differ between the two designs. Although such variations are monitored and compensated for, it is good practice to take advantage of cancellation. " Testing of this sort at the ATF requires, of course, that the antennas under test include one which will be selected for production.

Mechanical

Predictable mechanical effects arising from differences in antenna design, such as axis non-intersection, nodal point change due to tilt changes (gravity) and thermal deformations could be applied in real-time through a pointing model particular to each telescope design within the array. This procedure complicates an already complex instrument, and predictive removal of errors is unlikely to be as suitable as cancellation across an array of identical antennas. Mechanical decorrelation of the signal may also be difficult to compensate for antennas of differing design.

Appropriately chosen but different designs may provide cancellation of errors. Since pointing and beam shape variations arise from variations in the design of the primary and secondary mirrors, variations in the design of the quadrupod structure and details of the receiver design, a set of designs which minimize differences in these could mitigate image deterioration as discussed above.

Other

Lesser effects may also impact the scientific production of an array including antennas of different design. Webber notes that the arrangement of modules is likely to be somewhat different in different antennas. "In itself, this is only a minor nuisance. However, the air circulation will be different and the operating temperatures of various modules mounted in the racks and in the front end assembly are likely to change differently, both as a function of outside air temperature and as a function of elevation. I think this is likely to lead to phase drifts which are different between two antenna types. One may argue that, if the individual assemblies meet the phase drift spec, then all is well--but I think the better the various drifts match each other, the better the performance of the array." Software interfaces will undoubtedly differ between designs, increasing the size and complexity, hence the schedule and cost, of the software effort in ALMA.

ALMA science will also suffer in an array of heterogeneous antennas since commissioning and maintenance will both be made more complex, draining resources which might otherwise be directed to improvement of ALMA's scientific output. As Darrel Emerson has noted: "ALMA is pushing the state of the art in many areas, in particular in the phase stability needed at high frequencies in order to achieve high dynamic range. Having identical antennas helps a lot, because some interferometric errors will cancel out, and identical pointing errors can be allowed for in software. Even so, it's questionable whether we'll really achieve the very demanding requirements. Losing the advantage of some errors canceling out, because of the identical antennas, will make it that much harder to meet our requirements.”

An Explanation of How ALMA Will Correct for the Wide Field Polarization Beam, and Why Two Different Antenna Designs Will Not Make A Difference for Wide Field Polarization Imaging

M.A. Holdaway (NRAO/Tucson)
Oct 31 2005

Abstract: My personal view is that it would have been much better to have a single 12m antenna design and not two different 12m antenna designs in ALMA, and that when managment overrides the rank and file scientists, morale suffers. HOWEVER, I don't think having two different antenna designs will make wide field polarization imaging any more difficult. In order to do it for a single antenna design, we need to measure the pol leakage beam and correct for the pol leakage. This technique will work down to some level, leaving spurious ghost polarization which follows the total intensity emission at something like 0.1-0.5% of the total intensity emission on a pixel-by-pixel basis, and scattering flux from very bright sources like the synthesized beam's sidelobes due to time variable spurious polarization (due to the pol leakage beam rotating with the parallactic angle), potentially further spoiling our ability to see low level polarization on pixels with low level total intensity emission (ie, limiting the polarization across the field to something like 0.005% of the image peak). We will describe this method which we must use anyway even if we have only a single antenna design, and assert that it should work similarly well for two different antenna designs.

Several people have been making a big deal of wide field polarization imaging with ALMA, but it seems from their words that they do not actually understand wide field polarization imaging. It should not be considered a big deal. The key insight into the way it will work comes from Carilli & Holdaway (1992), and Holdaway (1992) extends this insight to an algorithm that should work for ALMA. Carilli & Holdaway were trying to image extended polarization with the VLA in the weak disk and halo of NCG 253, which is dominated by a central starburst of about 1 Jy flux. At L band, NGC 253 requires a three pointing mosaic. One pointing was right on the starburst, so the polarization calibration on the starburst was pretty good. However, the other two pointings found the central starburst right at the half-power point of the primary beam, which was a disaster.

The on-axis polarization calibration leaves the polarization leakage beam with a null in the middle, but fractional polarization leakages of over 5% are observed around the half power point. If you look at the half-power point of the polarization leakage beam, you see that the magnitude and orientation of the polarization leakage changes as you go around the half-power circle. IF the VLA were an equitorial mount, or if we only observed in snapshots, the polarization leakage beam would not rotate on the sky. IF all antennas had the exact same polarization beam, this would be a very simple problem: the spurious polarization in the polarization image would be equal to the total intensity image times the primary beam times the polarization leakage beam. In fact, the polarization leakage beams for the VLA antennas are fairly similar, and a simple image plane correction was used for the NVSS to correct polarization images down to something like 0.5%.

What about the case of long integrations where the polarization beam rotates against the sky, and what about the case where different antennas have different polarization beams? The following pages will show how we will address these issues.

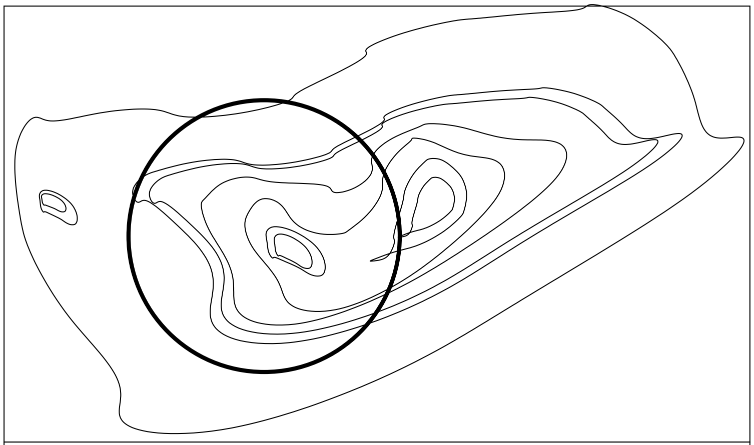

The top panel shows an extended source **I(x)** with a primary beam **B(x)** plopped down on top of it (where **x** is a vector coordinate, standing for (x,y)).

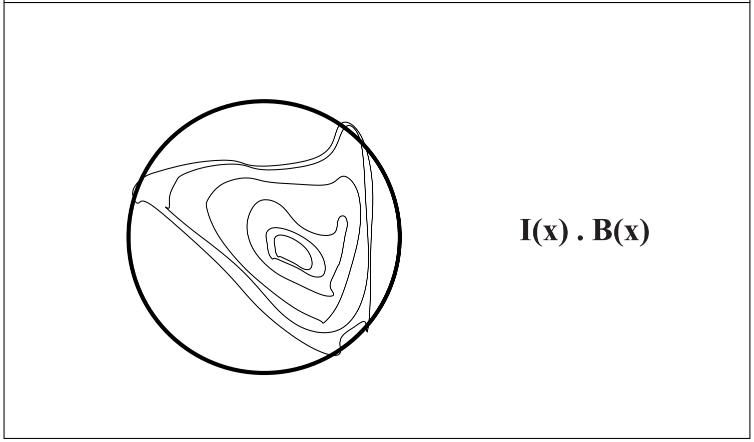

The second panel shows the extended source multiplied by the primary beam **I(x) . B(x)** . This is part of the observation process. The data visibilities are given by **FT(I(x) . B(x))**

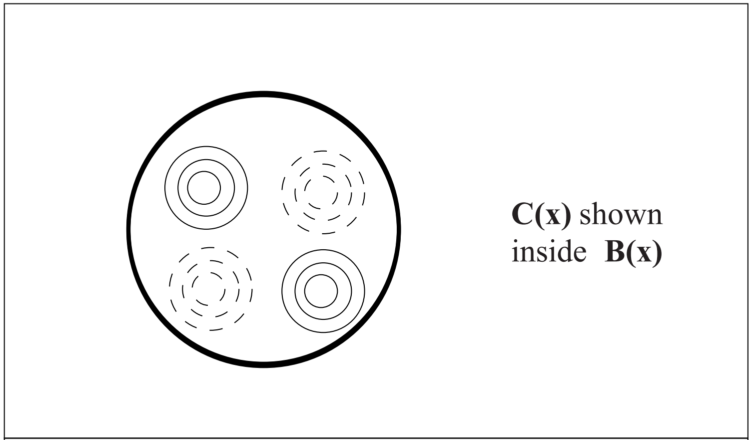

The polarization leakage beam **C(x)** for the VLA is a complex image with multiple lobes. The magnitude of the pol beam represents the apparent fractional polarization that an unpolarized source at that location in the beam would see. Each baseline will have a slightly different beam, though the average beam can be used for poor to moderately high quality images.

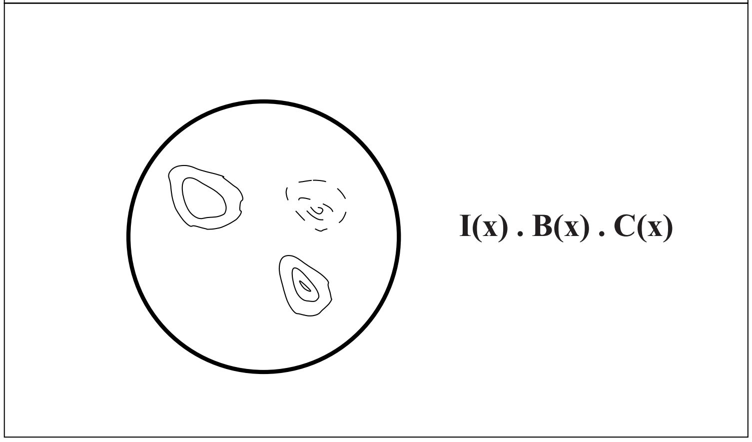

IF each antenna had the same pol leakage beam **C(x)** and we observed for just a snapshot, the spurious emission would just leak into the total intensity image as **I(x) . B(x) . C(x)**. To estimate the spurious pol emission for a single field, we would image and deconvolve **(I(x) . B(x))**, then multiply that image by **C(x).** To fix the pol image **P(x)**, **P'(x) = P(x) - (I(x) . B(x) . C(x))**. This exact correction was used in the NVSS snapshot images, and was accurate to about 0.5%.

Now consider a difficult case: we have significant extended emission so we require multiple pointings and total power observations; we observe over a wide range of parallactic angle so the simple image plane correction from the previous page does not work. However, we assume that all antennas have the exact same polarization beam (we'll treat the different polarization beam further on). This recipe, operating in both the image and visibility planes, should provide an excellent correction to the spurious polarization which leaks through the beam.

* Before observations start, we have done cross-polarization holography on an unpolarized point source and therefore have sampled the polarization leakage beams for the 12m ALMA antennas, the 7m ACA antennas, and the 12m ACA total power antennas. We have three pol beams, each an average for all antennas of the same kind.

* We do the regular on-axis polarization calibration along with phase and amp calibrations.

* We make the best total intensity image we can, using mosaicing and the ACA. (Note that our pol correction will only be as good as this image. If we have 10% errors in this image due to poorly utilized total power data, we won't do very well. Fortunately, ALMA will make high fidelity images.)

* Generally, ALMA will observe a different pointing on the sky every few seconds, though ALMA will often come back and reobserve each pointing several times at different parallactic angles. We must treat each snapshot observation separately. For each sky pointing/snapshot **i**, we will rotate ALMA's pol leakage beam **C(x)** to the correct orrientation for that snapshot's parallactic angle, and call it **C'_{i}(x)**. We then calculate the spurious polarization leakage in the image plane as:

I(x) . B(x) . C'_{i}(x)

However, rather than correcting for this, what we do is Fourier transform it and subtract this spurious leakage in the Fourier plane from the data visibilities, leaving residual visibilities which represent thermal noise and astronomical polarized signal. The residual polarization visibilities can then be mosaiced cleanly.

* In the case of mosaicing with the ACA, there will be three different pol leakage beams **C(x)**, so for each sky pointing/snapshot, there will be three different images **I(x) . B(x) . C'_{i}(x)**, so we will have to do three different Fourier Transforms (actually, only two: for total power, we just sum the image) to obtain the polarization leakage's contribution to the visibility data.

* In the case of Vertex and EIE antennas with different pol leakage beams **C(x)**, we will have to form five different images and Fourier transforms (ACA 7m, ACA 12m, Vertex **X** Vertex, EIE **X** EIE, and EIE **X** Vertex), and the Fourier transforms of these five different images will result in the pol leakage's contributions to the visibilities.

* In the case that every antenna has a significantly different leakage beam, we will have to calculate a leakage beam for each baseline, at which point finding the pol leakage's contribution to the visibilities will best be served by a single point DFT per baseline-image, rather than an FFT.

How well will this technique work? Such a calculation requires knowing how the image and beam errors are correlated, but we estimate. Assuming errors in **I(x)** are of order 2% (fidelity = 50, neglecting flux scale errors which will be the same for the pol image), and **C(x)** has max leakage of about 5% and is accurate to about 0.2%, we see the two contributions to the errors will be of order 0.1-0.2%.

Carilli & Holdaway, "A Simple Minded Approach to Polarization Mosaicing", VLA Test Memo 163, 1992.

Holdaway, "New Possible Algorithms to Improve the VLA's Polarization Performance", VLA Scientific Memo 163, 1992

Appendix. ASAC Report

In April 2003, the ASAC responded to its Charge from the ALMA Board, which inquired: "`Inhomogeneous Array: The ASAC is asked to discuss the impact on development, commissioning, and scientific performance of using two different 12m antenna designs in the ALMA array. Issues the ASAC is requested to consider include the impact of an inhomogeneous array on (i) the phase stability of the array, and its dynamic range and other imaging characteristics; (ii) operations and maintenance costs; (iii) software development schedule and costs; and any other issues the ASAC feels the Board should be aware of.`

From the Executive Summary of the ASAC report of the April 2003 meeting, presented to the ALMA Board at their 26-7 May meeting and noted in those minutes:
...
6. Inhomogeneous Array: The ASAC strongly recommends that a single antenna design be adopted for ALMA. Having two different antennas designs seems certain to impact the science capabilities of ALMA for wide field mosaics and polarization observations, while in a worst-case scenario, imaging of any significantly extended source could be affected. The ASAC recommends that the project consider whether additional specifications will [be] required to enforce consistency between two different designs. If two different antenna designs must be adopted, the ASAC recommends that an identical quadrupod design be used for both antennas, which should reduce any adverse effects on the science.

Later in the report...

The ASAC reviewed two written documents on the impact of an inhomogeneous array that had been prepared by A. Wootten and by the ANATAC. We also heard a presentation by S. Guilloteau. The science implications of having two different antenna designs arise primarily from "common mode errors", which would cancel if the antennas were identical. [This may also occur with the ALMA antennas built to identical specifications but differing in design details.] Common mode errors identified include pointing errors, phase/pathlength/focus errors, phase effects due to changes in the fiber length, and polarization matching and primary beam shape.

For common mode pointing errors, errors due to wind are likely to be common in the compact configuration, while solar heating in this configuration may vary from one antenna to the next due to shadowing. In contrast, in more extended configurations, common pointing errors are likely to arise from solar heating, while the wind and its associated pointing error may vary across the (large) site. For errors in phase due to pathlength and focus changes, all mechanical deformations except that due to the non-intersection of the axes (likely the dominant effect) would benefit from having identical antennas. Phase effects due to changes in the fiber length are dominated by the run to the antenna; this normally common mode error could probably be monitored and compensated

for in software. Polarization and primary beam shape are determined by the quadrupod leg design; having two different antennas with very similar quadripod designs could mitigate the problems here. However, it is worth noting that the Vertex and Alcatel prototypes do not have identical quadrupod designs. Inhomogeneous array designs also have cost implications during the construction, commissioning and operations phase. In the construction phase, the cost effect could be either positive or negative, depending on the details of the antenna contracts. For commissioning and operations, it is clear that having an inhomogeneous array implies extra costs due to the extra work involved with commissioning and maintaining two different antennas, maintaining two software interfaces (for example, different pointing models), etc. The bottom line is that anything that increases the cost ultimately affects the science return from ALMA in a negative way.
The ASAC reached the following conclusions concerning the inhomogeneous array:

1. The ASAC strongly recommends that a single antenna design be adopted for ALMA. Having a single antenna design will factilitate several key observing modes with ALMA, in particular polarization observations and wide-field mosaics. It will also reduce the effort and cost required to commission and operate ALMA.

2. If two different antenna designs are adopted, the ASAC recommends that the identical quadrupod design be used for both antennas. Having an identical quadrupod design should help to minimize science impact, again particularly for polarization and mosaic observations. Minimizing the problems introduced by having two different antenna designs implies that there should be additional specifications placed on the designs, for example on the lack of axis intersection, the thermal coefficient for the expansion of the quadrupod legs, the profile for the quadrupod legs, etc. It might be possible to minimize common mode errors with appropriate specifications on the change of the antenna with tempreature and gravity and on the wind response. However, placing a number of additional specifications on the antenna designs could drive the costs up.

3. If ALMA consists of an inhomogeneous array without stringent specifications on the quadrupod and other aspects of the 12m antennas, the ASAC believes the biggest potential impact on the science capabilities of ALMA will be in the areas of polarization observations and wide field mosaics. Polarization mosaics are probably the most demanding use of ALMA and would likely be extremely difficult with an inhomogeneous array. In a worst case scenario, imaging of any sources larger than roughly 1/4 of the ALMA primary beam could be adversely affected.

For any type of inhomogeneous array, the potential extra costs involved will take money and effort away from other ALMA tasks and the end result will be a less powerful instrument. Having two types of antennas will have a negative impact on commissioning and operations, with extra training, software, spare parts, etc. required. In this context the ASAC wishes to highlight the impact on the software effort, as many of the corrections required to operate ALMA with different antennas will fall to software.

Summary

The ASAC strongly recommends that a single antenna design be adopted for ALMA. The ASAC recommends that the project review the antenna specifications to see whether additional specifications would be required to enforce consistency between two different antenna designs. If two different designs are adopted, the ASAC recommends that the identical quadrupod design by used for both antennas, which should help to minimize the impact of the different designs on science. If two substantially different antenna designs are adopted, the biggest potential impact on the science capabilities of ALMA will be in the areas of polarization observations and wide field mosaics. In a worst case scenario, imaging of any sources larger than roughly 1/4 of the ALMA primary beam could be adversely affected.