ALMA Memo 283
Observing Efficiency of a Strawperson Zoom Array

John Conway
Onsala Space Observatory
S-43992, Sweden
email: jconway@oso.chalmers.se

January 5th, 2000

Abstract

We compare the observing efficiencies of variable resolution zoom arrays to conventional fixed array designs. To investigate this question in detail we analyse a specific strawperson zoom array. This strawperson array follows a logarithmic three-armed spiral on intermediate scales (radius 100m -1500m) but smoothly interfaces to a ring and a dense pack for the largest and smallest configurations respectively. Three telescopes are moved twice a week; using such a design it is feasible to have two full cycles per year through the $\leq$ 3km configurations.

We find that slightly fewer baseline-hours are lost to reconfiguration and calibration using zoom rather than fixed array designs. However, a much more significant advantage of zoom arrays is the ability to largely avoid data tapering in making images which require specific resolutions. A large fraction of ALMA observations which demand inter-comparisons of different line transitions or with optical and IR images will be of this type. We estimate that by avoiding tapering zoom arrays can observe approximately twice as many such sensitivity-limited fixed-resolution projects per year. We make a first attempt to consider the impact on zoom arrays of practical constraints such as wind delays on reconfiguration and fluctuations of demand. It appears that these do not effect the fundamental conclusions that zoom arrays can have higher efficiency than fixed arrays for a slightly lower construction and operating cost.

1 Introduction

Several recent ALMA memos have considered the question of how the number of configurations and their design effects the observing efficiency of the array. Observing efficiency measures the impact on the array productivity due to losing data to reconfiguration or tapering. Array efficiency is one important criteria in deciding between a small number of fixed configurations and a continously reconfigurable or 'zoom' array. In this memo we consider this question in detail. To make the discussion more concrete a 'strawperson' zoom array covering all the arrays sizes < 3km is proposed, and compared directly with fixed array designs.

2 A Strawperson Zoom Array

2.1 Spiral Path

In Figures 1 to 6 we show the geometry and operations of an example zoom array. The array is based on a three-armed spiral geometry (Conway 1998). At intermediate resolutions the spiral has a constant pitch angle. The spiral pitch is chosen so the azimuth changes by approximately 1.5 turns for a ratio of 8 in radius. At these intermediate scales the pattern is truly self-similar. A reduction in scale gives rise to exactly the same uv coverage (except for an unimportant rotation of the approximately circularly symmetric uv coverage). For the largest configurations the array is limited by the terrain to occupy an approximately circular plain about 3km in diameter. In order to allow for the maximum resolution given this terrain the largest configuration should approximate to a ring. This has been achieved by increasing the pitch of the spiral at large radius so that the spiral path smoothly joins a ring of radius 1.5km.

At the smallest scale the constant pitch spiral must be modified to take into account the minimum antenna-antenna separation that is allowed (15m between antenna centres is assumed). This is achieved by making the radius $r \propto \phi^{1.5}$ (where ϕ is the azimuth) in the inner parts rather than $r \propto exp(\phi)$ in the constant pitch part of the spiral. The packing density in the most compact configuration is then slightly less than the maximum achievable with hexagonal close packing of the antennas. Such maximum density hexagonal packing has several disadvantages including high baseline redundancy and sidelobes, shadowing losses and problems of access of the transporter to the antennas. In contrast the spiral design presented avoids most of these problems, giving slightly higher resolution and only slightly lower surface brightness sensitivity than a maximally packed hexagonal array. The details of this smaller array including its mosaicing performance have to be investigated, and somewhat higher density packing might be necessary. Again the overall strawperson design is not yet optimised for uv coverage or imaging but serves as a basis for discussing the array efficiency.

2.2 Pad Layout and Reconfiguration

Each arm of the array has a total of 60 pads, 19 pads in the $r \propto \phi^{1.5}$ portion, 30 in the constant pitch spiral portion and 11 on the ring. Going out from the centre these pads are numbered 1 to 60. The pads are randomly displaced from the spiral path both in azimuth and radius (with standard deviation 0.03 radians and 0.03 of radius respectively). Such random perturbations are necessary to break the symmetry and will be forced anyway by the need to avoid difficult terrain.

Placed on the pads are 21 antennas per arm, in addition there is 1 antenna at the centre of the array which never moves, giving a total of 64 antennas. The telescopes start by occupying pads 40-60 on each arm (see Fig 1), so that most of them are close to the 3km diameter ring. The resulting uv coverage is close to uniform. Two days per week, on Mondays and Thursdays there are antenna moves, on each of these days 3 antennas are moved, one on each arm. Each reconfiguration of 3 antennas can be accomplished by 3 transporters each moving one antenna taking approximately 1 hour for all < 3km ALMA configurations (Radford, 1999). These reconfigurations would be done before 9 local time when winds are at their lightest. Sharing some personnel between different transporters (Radford, 1999) a crew of 6 would be required, these personnel would be used for general array maintenance when not moving antennas. After moving the antennas there would be another 2 hours or so of antenna calibration before the moved antennas were ready to join the array. Usually during reconfigurations antennas would be moved directly from the highest numbered occupied pad to the highest numbered unoccupied pad on that arm; however in order to get more short spacings it may be advantageous to move directly to a lower numbered pad (such as pads 1 and 8 as shown in Figures 1 to 6). Clearly there is some scope for optimising the reconfiguration schedule.

With the above reconfiguration scheme on average 6 antennas/week, are moved until the most compact configuration is reached in which pads 1 - 21 are occupied. Every time there is an antenna move one of the pads 1-39 is occupied for the first time, since there are 2 antenna moves/arm/week the whole cycle takes just under 20 weeks. Figures 1 to 6 show snapshots of the array and uv coverage intervals throughout the cycle. Note that only a subset of the total number of configurations are shown. In total there are approximately 40 of these configurations, corresponding to the configurations after the Monday and Thursday moves on each of the 20 weeks. Given that the resolution ratio between the largest and smallest configurations is about 27 each of the 40 configurations will differ from the previous one by a factor of 1.09 in resolution.

2.3 The Overall Reconfiguration Cycle

Fig 7 shows a proposed overall cycling strategy for ALMA. Given that we can cycle through the < 3km zoom configurations in 20 weeks it is possible to have two such cycles per year allowing day and night observations at any given array size for all sources, and rapid return of data to PIs. It is somewhat better to have the zoom arrays start with the largest configuration and gradually reduce the array size. If a scheduled reconfiguration is missed, due for instance to high winds (see Section 5.4), the PI then receives slightly higher resolution observations than requested which can then be tapered to the required resolution.

The zoom reconfiguration periods are phased so that the two occasions when the array is in its most compact configuration are in early May and early November. These months give a good compromise between having good opacity and relatively low winds allowing high accuracy mosaicing observations. One advantage of a cycling scheme which is synchronous with the year rather than having an asynchronous cycle (e.g Radford 1999) is that configurations can be optimised to seasonal observing conditions. Because the 10km array will share few antenna pads with the $\leq$3km configurations and because of the large distances involved in moving antennas to this configuration, reconfiguration to and back from the 10km is expensive in lost observing time (Radford 1999). Given this it makes sense to have only one such 10km configuration per year. Putting this in early summer when winds are at their lightest and there is no snow is the best practical choice. Scheduling this array as early as possible from mid-November to mid-December avoids the high opacity conditions of January-March. The other long reconfiguration from the most compact array to 3km is timed to occur before May 15th before high winds start to make reconfiguration difficult. After this long reconfiguration the array stays stopped for 6 weeks in the 3km array for much of the period in which winds are highest in June and July.

With the scheme shown in Fig 7, 2 antenna moves/antenna are needed for each of the two zoom portions, a further 2 antenna moves/antenna are required to move to and back from the 10km array and 1 antenna move/antenna to move from the dense pack to the 3km array. The total number of antenna moves per year is therefore 7 per antenna, or 441 (assuming one of the 64 antennas never moves).

2.4 Scheduling

Given the 6 month cycle of the array described above, there would naturally be two proposal deadlines per year. Rather than requesting a particular array observers would request on their proposal forms instead the frequencies of the proposed observations and the required resolutions. PIs would also tick boxes to indicate whether exactly matched resolutions (to within 10%) were required or not. Any required constraints on observing conditions would also be specified. The proposals would be graded according to scientific merit. The highest rated observations would be scheduled for the appropriate array size to give the required resolution and frequency. Amongst the highest rated proposals those requiring exact resolution would be scheduled first with the excess capacity filled by experiments in which the exact resolution was less important (see Section 5.2). The date of the experiment (or range of dates for experiments requiring specific conditions) would be then be set assuming a preset reconfiguration schedule (i.e such as the one in Figure 7). Users would have access to the statistics of array size demand so that weaker proposals could maximise their chances of being scheduled by requesting resolution/frequency combinations that were requested less often.

It is not envisaged that reconfiguration schedule for the upcoming 6 months would be adapted in detail to the demand from proposals. Although there is sufficient transporter capacity to significantly vary the reconfiguration rate (see Section 5.3), this might entail staff scheduling problems. However what might happen after some years of operation is that the reconfiguration schedule is adapted to give more time in array sizes which have proven popular and less time in unpopular ones (see Section 5.3).

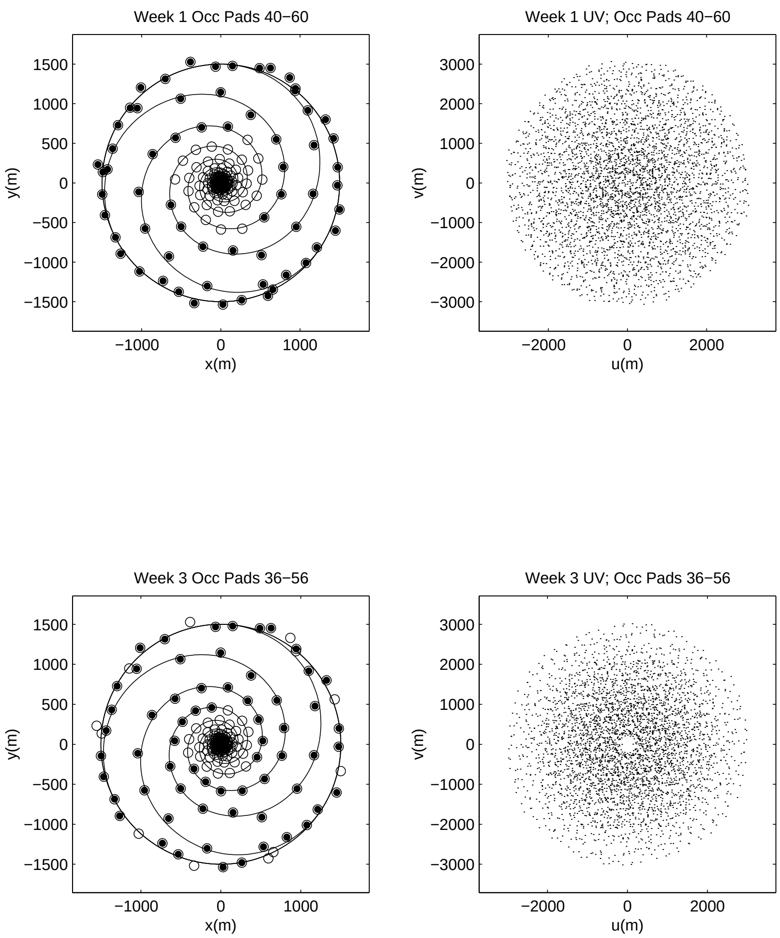

Figure 1: A selection of the 40 different configurations of the Strawperson zoom array introduced in Section 2. The black dots indicate the antennas and the open circles the unoccupied pads. The short spacing uv coverage of the week 3 configuration could be improved by forcing some occupied pads on the outer ring to lie closer together.

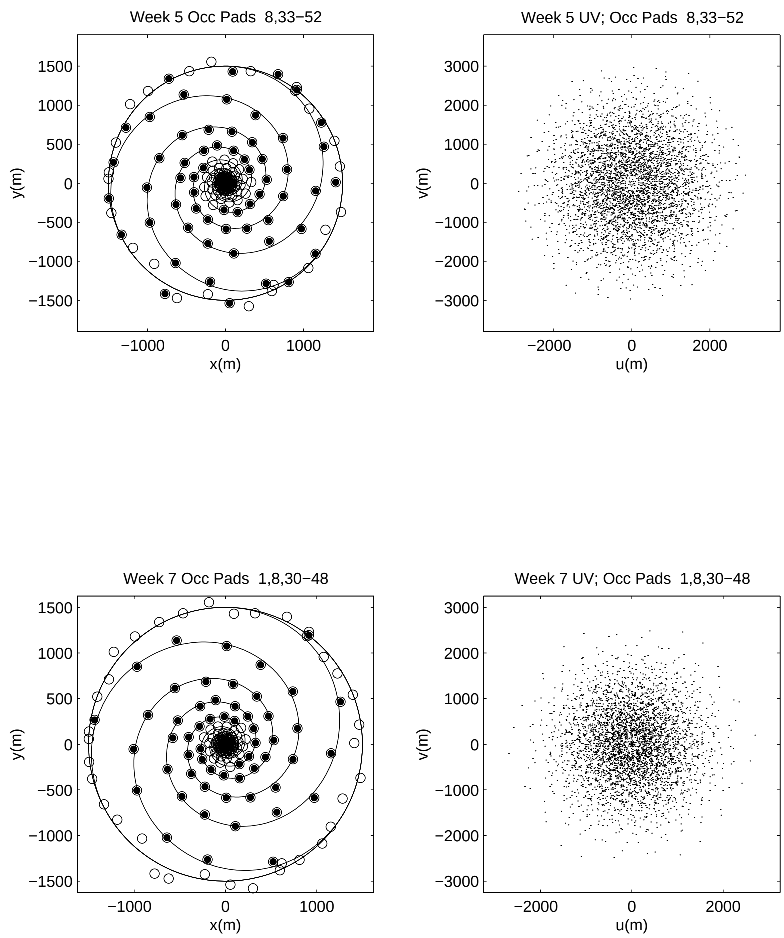

Figure 2: Between weeks 3 and 5, when moving from a ring like array to a spiral, one antenna is moved onto pad 8 of each spiral arm to give improved short spacing coverage.

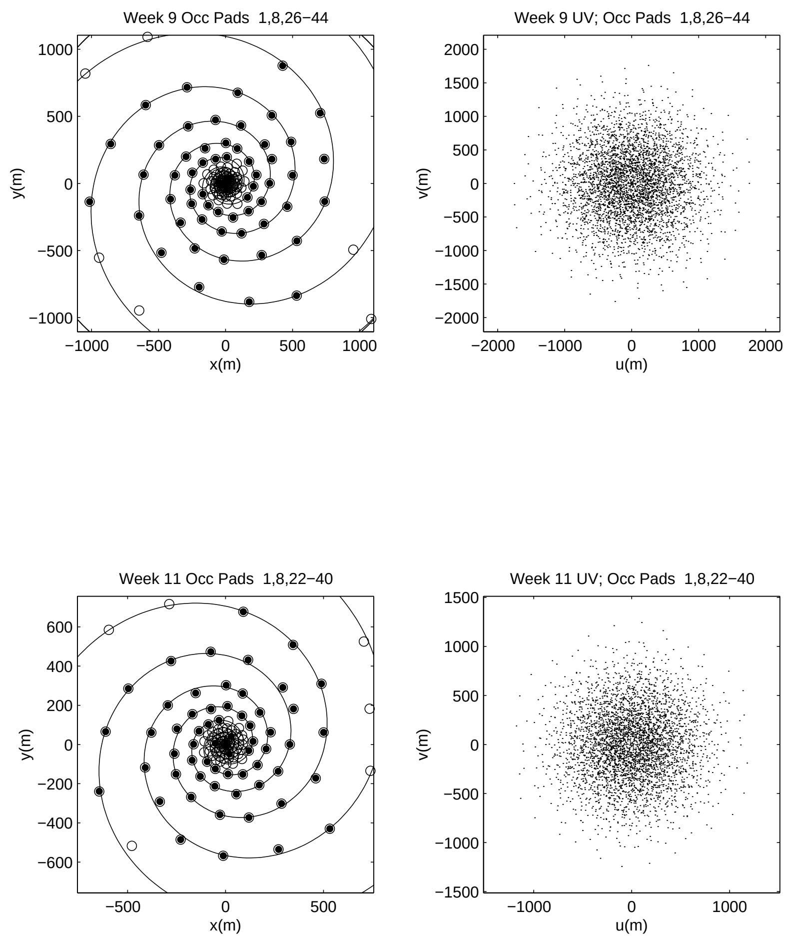

Figure 3: Between weeks 7 an 9 one additional antenna is moved onto pad 1 of each spiral arm to give improved short spacing coverage.

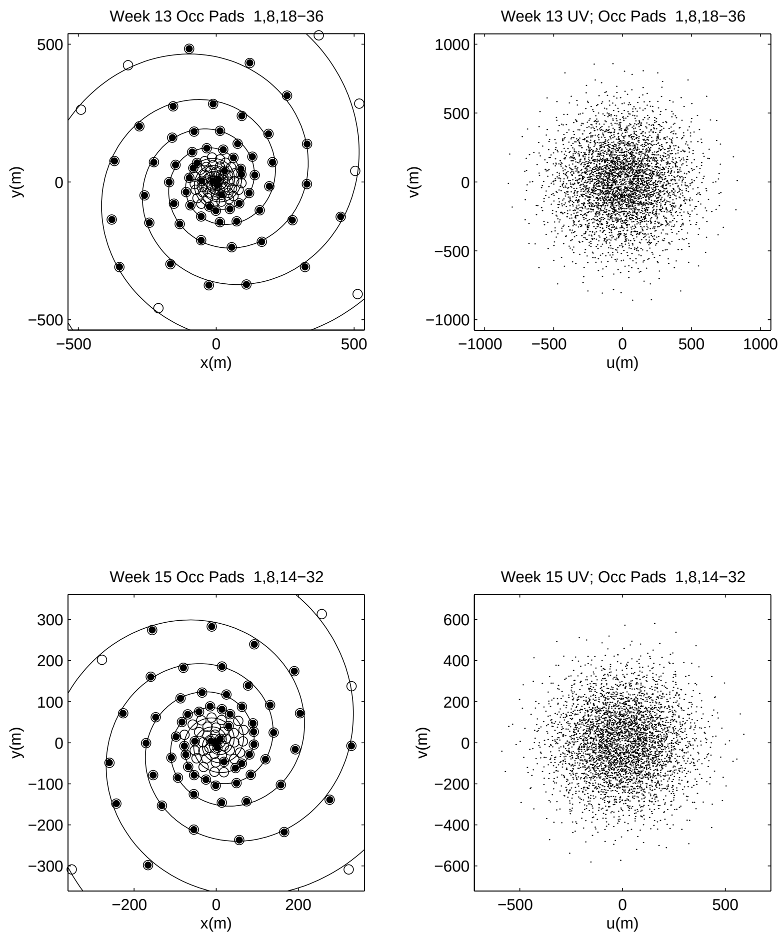
Week 13 Occ Pads 1,8,18−36
Week 13 UV; Occ Pads 1,8,18−36
Week 15 Occ Pads 1,8,14−32
Week 15 UV; Occ Pads 1,8,14−32
x(m)
y(m)
u(m)
v(m)

Figure 4:

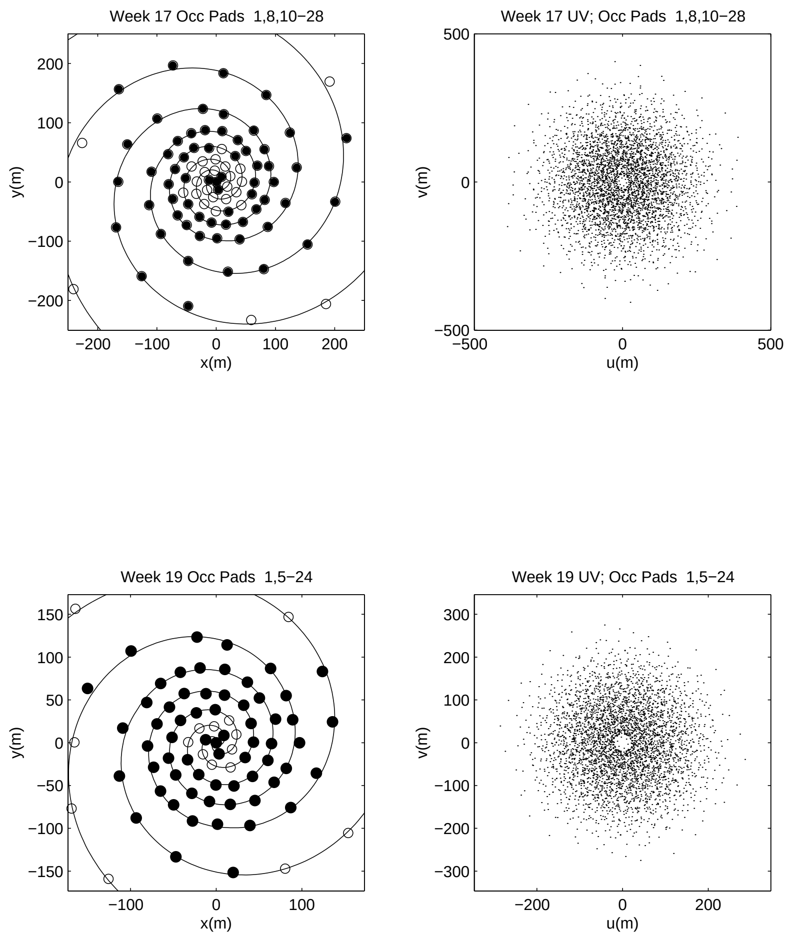
Week 17 Occ Pads 1,8,10−28
Week 17 UV; Occ Pads 1,8,10−28
Week 19 Occ Pads 1,5−24
Week 19 UV; Occ Pads 1,5−24
x(m)
y(m)
u(m)
v(m)

Figure 5:

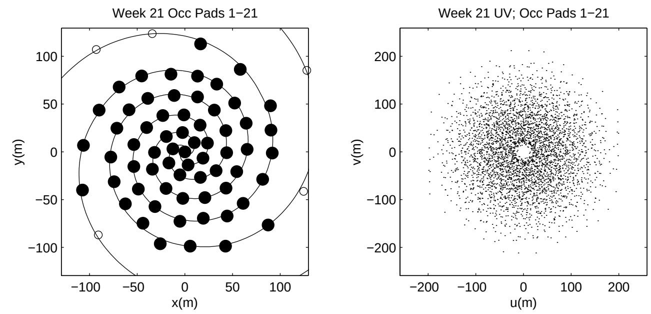

Figure 6: The most compact configuration. The size of the black dots approximately equals the antenna diameter of 12m. A slightly more compact packing could be achieved using slightly different array input parameters. The uv coverage shows one dot per baseline, in fact each antenna pair samples a circular region of 12m. The inner hole can be filled in by using mosaicing to recover the minimum baselines between the nearest edges of antennas.

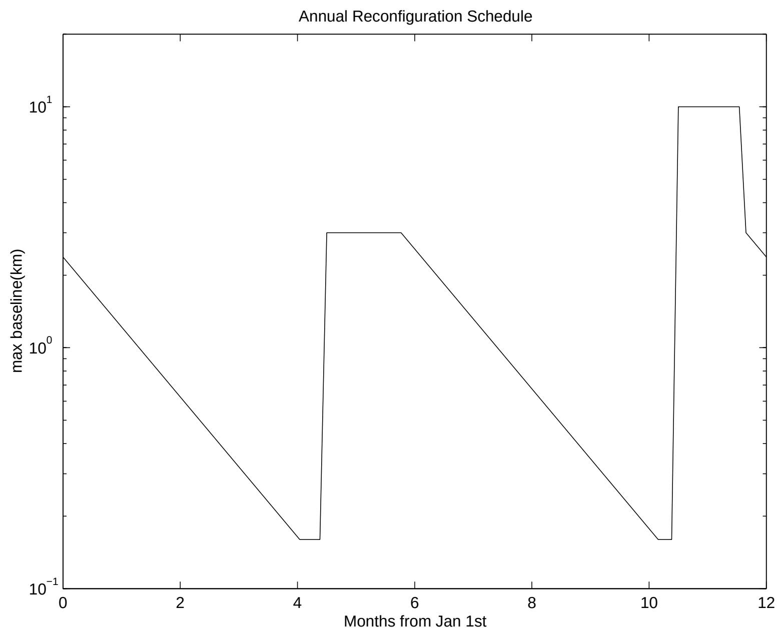

Figure 7: Suggested annual reconfiguration cycle. Two zoom cycles from 3km to 0.16km are included, and one period in the 10km configuration. There are three long reconfiguration per year, two to and from the 10km configuration and one from 0.16km to 3km.

3 Array Efficiency

Holdaway(1998) introduced the concept of total array efficiency, ϵ_{total}. Assuming a set number of experiments of equal length are observed per year, ϵ_{total} measures the loss in sensitivity for each experiment (sensitivity $\propto$1/rms map noise noise) compared to an ideal array which continuously observes with all telescopes and is never tapered. Alternatively one can say that the number of experiments that can be scheduled at fixed sensitivity scales as ϵ^2_{total}. Holdaway (1998) divided the total efficiency into two parts, 'Configuration Efficiency', ϵ_{conf} and 'Observing Efficiency' ϵ_{obs}, such that $\epsilon_{total} = \epsilon_{conf}\epsilon_{obs}$. Configuration efficiency measures the loss in sensitivity due to losing baseline-hours from the array when it is being reconfigured. The 'Observing Efficiency' instead includes the loss in sensitivity arising from the need to taper some significant fraction of ALMA observations to achieve a desired resolution.

In considering the effects of tapering one can usefully divide ALMA experiments into two types, those that do not require an exact resolution and those that do (Holdaway 1998). The former type of experiments, which here we call 'Type 1' or 'qualitative' science dominate on the present VLA. In these experiments one is after the overall structure of the source, and the data will not be compared in detail to other images. Experience on the VLA has shown that for these types of experiment obtaining a resolution within $\approx \sqrt{3}$ of the desired optimum is usually sufficient to achieve the science goals. Moreover, when the aim is to inventory all the scales of structure in a source at a given frequency multiple array observations factors of 3 apart in resolution are found to be sufficient. In contrast 'Type 2' or 'quantitative' science experiments require an exact resolution, and if the natural resolution of the array is not correct we are forced to taper the uv data; losing sensitivity. Such Type 2 experiments will be much more common on ALMA because it is primarily a spectral line instrument. A large fraction of the ALMA science is expected to come from making line ratio maps, requiring the same resolution at two transitions. Another such application will be comparing ALMA line images to images made from other instruments with their own fixed resolutions (e,g NGST and ground based optical/IR telescopes).

Executing Type 2 science programs with a small number of fixed arrays requires very extensive tapering of the data (Holdaway 1998). Since tapering effectively deletes the long baseline data which we have gone to such trouble to collect, the result of this is a significant loss in array sensitivity and hence observing efficiency ϵ_{obs}. Yun and Kogan (1999) have analysed the effect on ϵ_{obs} of tapering of data for set array configurations as a function of the number of configurations, N_c. The exact result depends on many hard to predict parameters of array use. However assuming 50% of observations are of Type 2 they find for $N_c = 4$ that the overall $\epsilon_{obs} \approx 0.85$. Considering only Type 2 projects however ϵ_{obs} is between 0.6 and 0.7.

Next we turn to consider the efficiency of zoom arrays, first considering the configuration efficiency and then the observing efficiency.

4 Zoom Array Configuration Efficiency

Two recent memos (Guilloteau(1999) and Yun(1999)) have considered the loss of observing time that results if one reconfigures rapidly (as many antennas as possible every day, before re-calibration) or slowly, moving only one or two antennas per transporter per day before recalibration. The for-

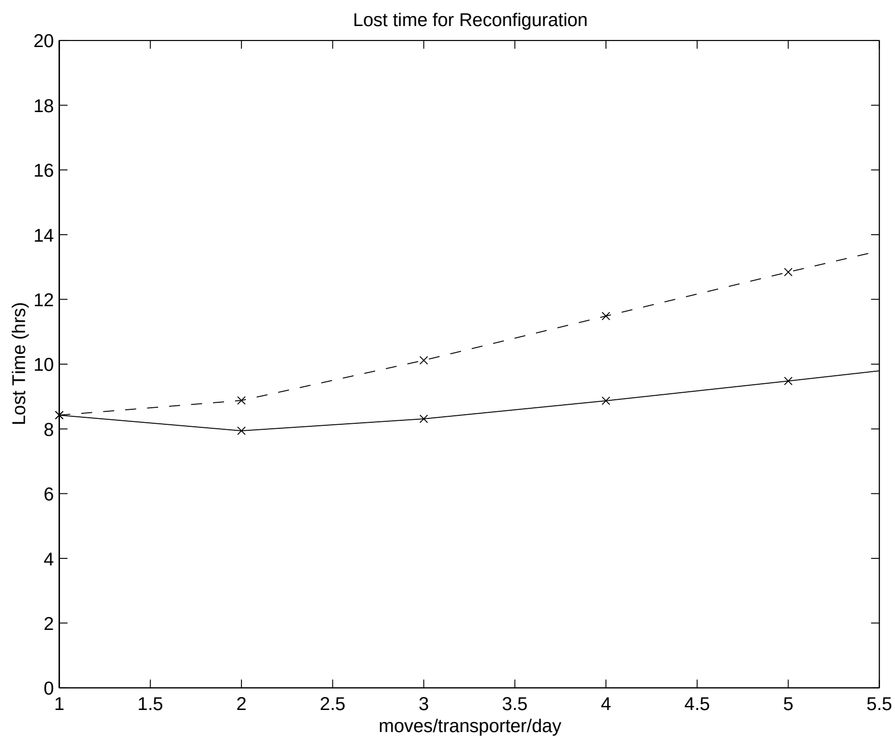

Figure 8: Estimate of equivalent array lost observing time, τ_{lost}, accrued after moving all the antennas once, as a function the number of antenna moves/transporter/day, (N_m). Assumes $N_t = 3$ transporters, antenna move time $T_{move} = 1$hr and calibration time $T_{cal} = 2$hrs. $N_{extra} = 2$ unmoved antennas are also assumed required during calibration, The dotted line uses the movement model discussed by Guilloteau(1999) and Yun(1999). The solid line uses a more realistic model discussed in the Appendix.

mer is the expected reconfiguration strategy for set configurations and the latter for zoom arrays. Somewhat counter-intuitively both memos found that slow reconfiguration loses less observing time than rapid reconfiguration. The Appendix discusses the origin of this result and the effects of taking more realistic assumptions about move strategy and move time. The effect of these modifications, shown in Figure 8 (solid curve) is that if the total number of antennas moved per day is less than about 15 there is in practice little difference in efficiency between moving fast or slow.

The quantity 'lost observing time', τ_{lost} in Figure 8 is in fact defined (Guilloteau 1999), as the extra time the whole array would need to observe to make up for the baseline-hours lost due to a complete 'reconfiguration' of the whole array. A better way of stating τ_{lost}, more applicable to true zoom arrays, is that it is the lost observing time accrued by the time *all the antennas have been moved once*. The total lost time in moving through a set of configurations depends therefore only on τ_{lost} times the number of moves per antennas to cycle through all the configurations. It follows that zoom arrays have a slight advantage over set arrays not because τ_{lost} for slow reconfiguration is significantly less than for fast reconfigurations but because in general fewer moves/antenna are required to sample all arrays.

For our strawperson array array discussed in Section 2 the total time lost to reconfiguration is 8.5 hours x 2 moves/antenna = 17 hours. Since the cycle takes 20 weeks the fractional observing time lost to reconfiguration is very small, 0.5%, and the resulting $\epsilon_{conf} = 0.9975$. This configuration efficiency for our 40 configuration zoom array differs radically that that obtained by Yun and Kogan (1999) who found that as the number of configurations increased the lost time increased linearly and became very significant for $N_c > 8$. It was this effect that limited the maximum number of configurations to between 4 and 6. The main reason for this very different conclusion was that Yun and Kogan assumed that no pads were shared between configurations so that *all antennas were moved every time there was a reconfiguration*. While a reasonable approximation for the case of a small number of configurations, this approximation completely breaks down when N_c is large and one can arrange that different configurations share most antennas. For example for our strawperson zoom array of 40 configurations only 3 antennas are moved between each configuration and not 64.

5 Zoom Array Observing Efficiency

5.1 Ideal Case

In estimating the observing efficiency of a zoom array we first consider the ideal case in which the demand on the configuration maximum size (B_{max}) is assumed perfectly flat as a function of Log B_{max}; and so exactly matches the 'supply' function of a zoom array reconfiguring at a constant number of antennas per week. All projects are then assumed to be scheduled to be observed in the configuration which gives the natural un-tapered resolution as close as possible to that desired. Since the scale factor between configurations is only 1.09, each experiment will be observed with a natural resolution within 1.09 of that desired. Experiments which require exact resolution will then be tapered. Figure 9 shows the increase in noise that occurs when the resolution is reduced by tapering. For the maximum tapering required the noise increases by $\approx$1% implying $\epsilon_{obs} > 0.99$.

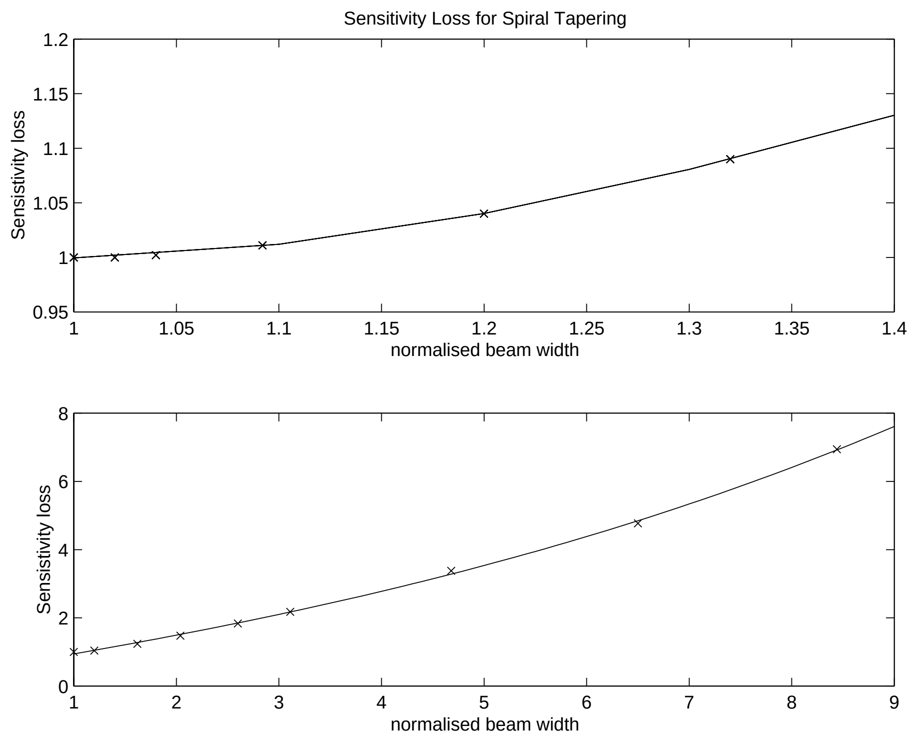

Figure 9: Sensitivity loss due to tapering when the array is in its intermediate logarithmic spiral configurations.

5.2 Statistical Variations

In reality of course, given the finite number of observations, the distribution of demanded maximum baseline sizes will not exactly match the supply, even if *statistically* the demand function is flat with Log(B_{max}). What is the effect of these statistical fluctuations on ϵ_{obs}? To evaluate this we assume for simplicity that all experiments have the same length (say 4 hours) and that each must be scheduled in its correct LST 'slot'. Consider the experiments competing for a particular slot (say 0 - 4 LST). For our strawperson array with a cycle time of 140 days there are of course 140 slots available at this LST range. We assume that the 140 highest graded proposals in this LST range are chosen to be observed. We will further assume that 70 of these are 'Type 1' and 70 are 'Type 2' observations (see Section 3 for definitions). Since the Type 2 observations require particular resolutions (see Section 3), these observations should obviously be scheduled first. Assuming a flat demand function as a function of $LogB_{max}$ then the optimum arrays requested by the 70 Type 2 experiments will be randomly distrubuted over the 40 possible. A typical random distribution of the demanded arrays for the 70 Type 2 observations is shown in Fig 10,top.

For our strawperson array plan Monday-Thursday configurations are 3 days long and those from Thursday-Monday 4 days long, so half of the configurations have 3 opportunities to schedule experiments in a particular LST slot and the other half 4 opportunities (see Fig 10, middle). If there are more Type 2 experiments demanding a certain configuration than supply available then some such experiments must be moved to a slightly non-optimum array. Various algorithms are possible for rescheduling the excess at a given configuration, here we simply assume they are scheduled one configuration earlier than requested. By observing earlier the resolution will be slightly larger than that desired and the exact resolution required can be recovered by gentle tapering. If after rescheduling the capacity for the new configuration is also exceeded one of the original experiments is moved on one configuration. The chosen algorithm therefore favours a larger number of small moves in experiment scheduling rather than a few large moves. In the example given in Fig 10, which is typical, only 4 Type 2 experiments out of 70 must be moved by one configuration from optimum. No experiment need be shifted by more than one configuration. The final distribution of scheduled Type 2 experiments is shown in Fig 10, bottom. The remaining capacity at each configuration is then filled up by scheduling the Type 1 observations. For these observations achieving the exact resolution requested is not so critical, despite this even most Type 1 observations are also accommodated within one or two configurations of the resolution they would most desire.

We have carried out Monte-Carlo simulations which show that if 50% of experiments are Type 2 then for our strawperson arrays on average 6.5% of these experiments must be shifted by 1 configuration while only $< 0.1\%$ must be shifted 2 configurations. A shift of two configurations is so rare because such a shift is only needed if the excess of supply over demand in a given configuration is more than the *total* supply for the next largest configuration. If shifted by one configuration the resolution is at maximum only 1.18 times optimum. From Fig 9, the sensitivity loss on tapering such observations to the optimum resolution is only 3%. We also find that even Type 1 observations are virtually always scheduled within 2-3 configurations of their best preference array, which is better than for the case of four set arrays separated by factors of 3 in resolution.

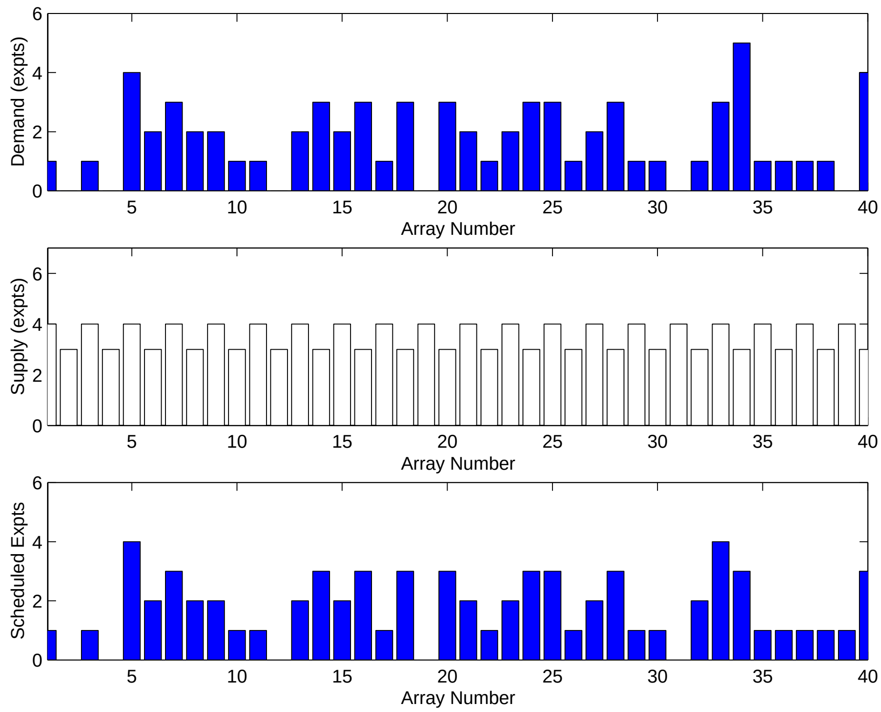

Figure 10: Effect of statistical fluctuations in the distribution of requested array number for Type 2 observations (see Section 5.2). Top - a typical distribution of demanded configurations, assuming a 140 day cycle with 140 LST slots and 70 type 2 experiments randomly distributed over the 40 possible arrays (i.e. according to a flat probability distribution as a function of $Log(B_{max})$. Middle - the number of LST slots available in each of the 40 arrays, Monday-Thursday configurations have 3 available LST slots and Thursday-Monday 4 available slots. Bottom - The distribution of Type 2 experiments after moving experiments which exceed the capacity to an earlier array. A total of 4 experiments had to be shifted by one configuration, which is typical. The remaining supply at each configuration will be filled up with the 70 Type 1 observations.

5.3 Non-Flat Demand Function

In keeping with Holdaway(1998) and Yun and Kogan(1999) we have in Fig 10 assumed a statistically flat demand function as a function of Log B_{max}; this to a first approximation correct for the VLA (Holdaway 1998). Is this a valid assumption? The question is somewhat complex because the number of proposals submitted for a given array size is itself probably closely coupled to the amount of time offered. Given the wide ranges of angular resolution demanded by ALMA science and its wide range of frequencies, the probability is that if we offer equal time in all arrays as function of Log B_{max}, the demand function for ALMA will also be flat. What however are the consequences for a zoom array if it is not? We can of course vary the rate of reconfigurations to match the demand- but what if we keep with a constant reconfiguration rate of 6 antennas/week?

Consider if the demand function varied by a factor of 2 over the cycle and was large (e.g 4/3 x 3.5 expts/LST slot/configuration) at intermediate sized configurations and small (e.g 2/3 x 3.5 expts/LST slot/configuration) for the largest and smallest configurations. We assume 50% of all experiments are Type 2 and this ratio is the same for all array sizes. For the peak demand for intermediate arrays, the demand for Type 2 experiments would be (e.g 2/3 x 3.5 expts/LST slot/configuration), which is still less than the average supply of 3.5 expts/LST slot/configuration. Monte-Carlo simulations show that in this case even at peak demand only 15% of Type 2 observations must be shifted by 1 configuration. Half of the remaining Type 1 observations could also be scheduled within a couple of configurations of their most desired configuration. The remaining half of the Type 1 experiments would however have to be displaced by about 10 configurations to find sufficient free LST slots, or a factor of 2.3 in resolution. In comparison for four fixed arrays - all experiments, irrespective of type, can be observed at a natural resolution of up to $\sqrt{3} = 1.7$ different from their most desired resolution. In this case the size of the resolution mismatch simply depends on chance; on whether or not the combination of frequency and the offered array sizes matches the resolution desired for the experiment. Although here the largest possible resolution mismatch can be larger for zoom arrays, these have the advantage that we can decide which of the experiments will be effected by the large resolution mismatches, and so we can reduce the impact on scientific return.

For the above case study an obvious way for the zoom array to perform better would be to vary the rate of reconfiguration. The simplest way to do this would be to have occasional holidays in reconfiguration in which no moves are made for a week or two, the overall configuration schedule allows time for such breaks (see Section 2). A single break near the configurations with maximum demand would accommodate most of the losses in the test case we considered above. A zoom array obviously has much more flexibility in this regard than a set of fixed arrays in which only the total time in each fixed resolution array can be varied. Conversely in a zoom array if less time needs to be spent in unpopular configurations the move rate can be temporarily increased. There is plenty of capacity for such speedup because the maximum move rate of between 4 and 5 antennas/transporter/per move day (Radford 1999) in all seasons except Winter (when winds might restrict reconfiguration) is much greater than the average required rate of 1 antenna/transporters/per move day. Configuration holidays and faster configuration could be scheduled in detail 6 months in advance to match demand once the proposals were rated. More likely however is that ALMA would start with zoom arrays reconfiguring at constant rates and based on experience would adapt to a standard reconfiguration scheme which best fitted demand averaged over several years.

5.4 Wind Delays

An important practical question is how often a reconfiguration will be delayed by high winds. The maximum wind speed at which antennas can be reconfigured is 16m/s, which is close to the median wind speed during winter afternoons (Radford 1999). A reconfiguration delay will mean that a project is observed at a non-optimum resolution. If the zoom array decreases in size with time as indicated in Fig 7 then the observed resolution will be higher than desired and quantitative Type 2 science projects will need to be tapered to get the desired resolution, causing decreased observing efficiency.

Wind speed is a very strong function of local time of day (see Figure 1, Radford(1999) and Figure 4, Holdaway(1996)) gradually increasing after sunrise to a maximum at 3pm local time. The effects of wind have been minimised in our chosen reconfiguration strategy (section 3) by using 3 transporters and reconfiguring on average only 1 antenna/transporter/move day, so that reconfiguration is completed before winds get high. If the work day starts at 8am local time, antenna moves will be completed by 9am and antenna pointing calibration by 10am. Fig 11 shows the fraction of the time winds are below the critical value of 16m/s as a function of time of day and season. The windiest months are June and July and for these months winds are above the critical value at 9am local time between 10% and 15% of the time.

If we are in the zoom portion of the schedule during the winter months we therefore expect that between 10% -15% of the scheduled reconfigurations will be missed. These missed reconfigurations could simply be rescheduled for the next day; alternatively we could simply move 2 antennas/transport on the next scheduled move day instead of the normal 1 antenna/transporter. Even if reconfiguration is delayed by 3 days corresponding to a whole configuration; then the effected experiments would be observed on average at 1.09 times their optimum resolution and the sensitivity loss for tapering is only a few percent. We would have to miss out 3.5 configurations - equivalent to having no reconfiguration for 12 days, before the resolution mismatch was as bad as for the case of 4 fixed arrays which scale in size by a factor of 3 (i.e. resolution mismatch of $3^{0.25} = 1.31$ on average).

Since wind is only a significant problem for two or three months a year another way to minimise its effects would be to phase the configuration cycle so that the array was stationary during these months. We could for instance schedule a long period in the 10km, 3km or 0.16km configurations in the winter. Fig 7 for instance shows the array stationary in its 3km configuration for much of June and July.

Clearly more information on the statistics of winds at the site and the effects on the reconfiguration schedule are required. One important question is that of the wind constraints on the pointing calibration of the antennas. Despite this, it appears that by confining the reconfigurations to early in the morning wind delays should not be a dominant effect even in winter; and we might further reduce its effects by scheduling a static array for the windiest months.

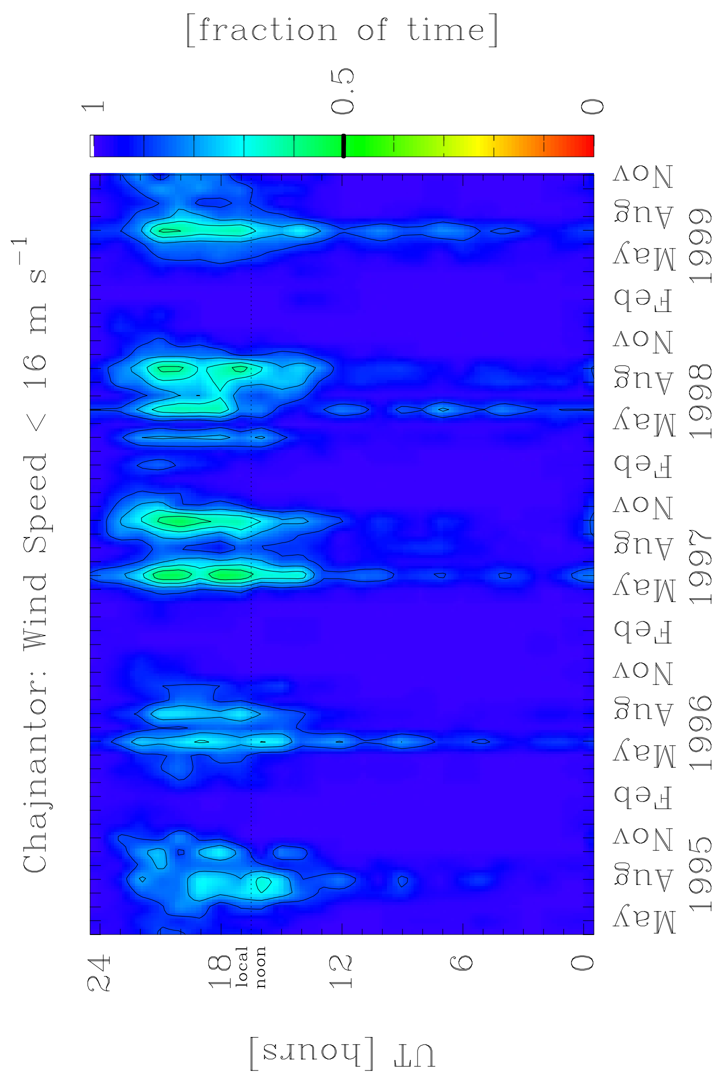

Figure 11: Plot showing the fraction of time wind speeds are below the antenna move limit of 16m/s as a function of time of day and time of year. Chilean civil time is UT-4hrs in the southern hemisphere winter and UT-3hrs in summer. Data taken from the site data archive (at http://www.tuc.nrao.edu/alma/site/Chajnantor/overall/threshold-windspeed-16.ps), S.Radford, private communication. Contours are at 10% intervals.

5.5 Dynamic Scheduling

A significant fraction of ALMA projects will require particular observing conditions of opacity and phase-stability; it is therefore likely that ALMA will used 'dynamic scheduling' where experiments are scheduled provisional on certain observing conditions holding. Some of these experiments will also be of Type 2 and require exact resolution. On a zoom array, if the delay is likely be less than a day or two, to maximise the chances of success these experiments should obviously be scheduled on the first day of the 3.5 days of the most appropriate configuration. If the required weather conditions do not occur within that configuration then the observations will have to be delayed and made in a later non-optimum configuration, losing observing efficiency.

The situation can be compared favourably to fixed arrays a factor of 3 apart in resolution; in which as noted in Section 5.4 on average the resolution achievable at any given frequency is within a factor 1.31 (the fourth root of 3) of that desired. Analogous to the case of wind delays a dynamic observation on a zoom array would have to be delayed by on average 3.5 configurations or 12 days to give a resolution inaccuracy as bad as this. The situation for zoom arrays can be further improved by scheduling the first opportunity to observe a dynamic project in a slightly larger array than the optimum given the desired resolution; in the anticipation of an observing delay. Doing this zoom arrays should outperform fixed arrays provided the average delay was of approximately < 20 days. This should encompass most observations. We note that the most demanding sub-millimeter observations would probably in any case be made when the array is spending a few weeks stopped in its compact configuration at the end of a zoom cycle (see Fig 7).

6 Conclusions

Table 1 summarises the main parameters of our strawperson zoom array compared to an array with 4 set configurations, for the < 3km arrays of ALMA. The '4-fixed' array is assumed to share no pads between configurations. In contrast the '4-fixed-shared' array is assumed to share 30% of the pads between configurations (extrapolating from the 36 antenna design of Kogan(1998b)). Table 1 shows that the zoom array has fewer pads and requires fewer antenna moves to go through the configurations than the fixed array designs. The overall operating cost is almost linearly related to the number of antenna moves (see Yun and Kogan(1998)) and is therefore also slightly less for the zoom array than for the fixed arrays.

6.1 Comparative Efficiency

The configuration efficiency of the zoom array in Table 1 is calculated as described in the Appendix, following Guilloteau(1999), by calculating the number of baseline-hours of observing time lost due to telescope motion and calibration. For the fixed arrays an upper limit to the configuration efficiency is calculated in the same way, assuming 16 antennas are moved and calibrated every move day, and all available baseline-hours are used for scientific observations. A lower limit to the efficiency is set instead using the model of Yun and Kogan(1999), and assuming that none of data from the reconfiguration days is useful. This would be the case if there were little demand for the poor uv coverage of a hybrid array or because mixing rapid antenna reconfiguration and recalibration of the array was found to be impractical. Our conclusions are that the zoom arrays

Array Type	Pads	Antenna Moves	No of Arrays	Resolution step	ϵ_{conf}	Overall ϵ_{obs}	Type 2 ϵ_{obs}
4-Fixed	256	192	4	3	0.953-0.996	0.85	0.65
4-Fixed-Shared	195	134	4	3	0.967-0.997	0.85	0.65
Zoom	181	117	40	1.09	0.9975	0.99	0.98

Table 1: Comparison of a Strawperson Zoom and Fixed Arrays for the < 3km sized configurations of ALMA, assuming 64 antennas. The 4-fixed array shares no pads between configurations while the 4-Fixed-Shared array shares 30% of pads between configurations. In all cases configuration efficiencies are calculated assuming we move from the largest to smallest configuration in 20 weeks. The calculations for efficiency and total number of antenna moves therefore do not include moving to and from the 10km array or back to the smallest array. Two figures are given for ϵ_{obs} for the fixed arrays, a lower and an upper limit, see Section 6.1 for details.

have higher configuration efficiencies than the fixed arrays; however as noted by Yun(1999) these efficiencies are relatively large for all arrays.

In contrast we do find a significant difference between zoom and fixed arrays in terms of observing efficiencies. The figures for the fixed arrays are taken from Yun and Kogan(1998). The observing efficiencies for the zoom array are lower limits assuming a maximum shift of 1 configuration for Type 2 observations (see Section 5.2), and using the tapering loss function given in Figure 9. It can be argued that the performance of zoom arrays relative to fixed arrays is likely to be even better than indicated in Table 1. The reason is that the fixed array calculation of Yun and Kogan(1998) assumes an array size demand function for Type 2 experiments which peaks at the array sizes offered and scales as $1/\theta$; in contrast the zoom array calculation assumes a completely flat demand function as a function of array size. If the efficiency loss for fixed array designs was calculated with this same flat demand function then the values of ϵ_{obs} would be significantly smaller than in Table 1. Even so Table 1 shows that the efficiency difference between zoom and fixed arrays is very significant especially for Type 2 experiments. This is even more striking if we consider the number of sensitivity limited Type 2 experiments that can be observed per unit time which is proportional to ${\epsilon_{obs}}^2$.

The main advantage of the zoom array is clearly that it allows very finely varying resolution which almost completely eliminates the need for the tapering of data. The zoom array therefore is much better in terms of the observing efficiency, especially for quantitative science (or Class 2 experiments). Overall the strawperson zoom array is slightly lower in capital and operating cost but has higher efficiency than the set arrays.

6.2 Variants

The zoom array design presented in this memo is only one of many possible. Adding pads would obviously increase performance for increased capital and operating cost. In addition there are alternative ways to run a zoom array, in a sense the design of a set of a self-similar pads with maximal sharing between configurations can be decoupled from how the array is run; truly continuous or in set configurations. For instance one could build a set of pads as shown in Figure 1 - 6 but operate it instead as set of fixed arrays a factor of 2 apart in resolution reconfiguring approximately one third of the antennas every month (Conway 1998). During a reconfiguration the uv coverage would gradually change and those experiments requiring exact resolution could be observed in these hybrid arrays. Such a mode of operation might make sense if the number of experiments requiring exact resolution are in a minority. Once however the number of such projects approaches 50% a true continuously evolving array is likely to be more efficient.

6.3 Imaging and Site Constraints

This memo has primarily considered the question of zoom array observing efficiency. This of course is only one aspect of deciding between zoom arrays and conventional fixed arrays. The other important questions are the imaging quality of zoom arrays compared to fixed arrays and whether a zoom array can be fitted into the terrain. Simulations with test images (Conway and Jerkstrand, Memo in prep) suggest that zoom arrays give similar imaging performance to other designs of arrays, as long as one ensures sufficient short spacings. These short spacings can be obtained either by placing telescopes on the innermost pads (see Figures 3,4,5) or by closely pairing together some pads on the spiral arms or outer ring (see Figure 1). Such short spacings in inhomogeneous array designs could also come from an array of small dishes in a compact configuration.

On the question of the terrain it seems that it is possible to fit an outer ring into the 3km diameter plain, and that there are also places where a $<$ 1km diameter spiral can easily laid out (Radford 1999). From Figure 1 it can be seen that about 20 pads lie in the difficult terrain between the ring and the $<$ 1km diameter region of the array. Since the spiral pattern can be rotated by an arbitrary amount with no effect on the uv coverage and the pads in any case require small perturbations from the spiral path to break the symmetry, it is probable that appropriate pad locations can be found (see Kogan (1998b)). Another important terrain constraint is that the $<$ 1km diameter region of the array should lie close to the centre of the 3km ring; such a 'concentric' constraint is of course a general property of most of the array designs proposed so far. Clearly further work is urgently needed to consider how in detail a zoom array design fits into the terrain.

Appendix: Configuration Efficiency

Guilloteau (1999) and Yun(1999) have both discussed the lost observing time that occurs due to reconfiguration as a function of the number of telescopes moved per day per transporter. Both find the apparently counterintuitive result that less observing time is lost if antennas are moved slowly than if they are moved fast. The origin of the effect can be seen if we consider in detail the baseline-hours of observations lost during reconfiguration.

Consider an array of N antennas, using N_t transporters and moving N_m antennas per transporter per day. Guilloteau (1999) and Yun(1999) both assume we take the whole set of moving antennas $N_{out} = N_m N_t$ out of the array for the length of time required to move all these antennas, $N_m T_{mov}$. Additional time T_{cal} is lost in recalibration of antenna pointing and positions. First let us consider the contribution to the lost baseline-hours due to antenna motion. The number of baselines lost to the array while moving is $N_{out}(2N - N_{out} - 1)/2$ which for $N_{out} << N$ is approximately $N_{out}N = N_m N_t N$. The number of baseline-hours lost per reconfiguration day is therefore approximately $N_m^2 N_t N T_{mov}$. Multiplying by the number of reconfiguration days required for a 'complete reconfiguration', or better stated, the number of days required to move every antenna once, the total number of lost baseline-hours is $\approx N_m N^2 T_{mov}$. The equivalent 'Lost Time' (τ_{move}) needed to observe to make up the sensitivity lost (Guilloteau 1999) equals the lost baseline-hours divided by the number of baselines in the full array, i.e. $\tau_{move} \approx 2N_m T_{mov}$.

The assumption that all antennas to be moved are taken out of the array for the length of time it takes to move *all* of these antennas, made by both Guilloteau (1999) and Yun(1999) can be challenged. More realistically if there were one transporter and it moved 3 antennas per day and each move took 1 hour, then for the first hour we have an array of 63 antennas, the second hour an array of 62 antennas and the third hour an array of 61 antennas. The formulations of Guilloteau (1999) and Yun(1999) instead implicitly assume that we have an array of 61 antennas for 3 hours. The result of correcting for this is that the number of baseline-hours or time lost due to antenna motion changes by a factor $(N_m + 1)/(2N_m)$, which approaches 0.5 as N_m becomes large.

If we consider the lost baseline-hours lost due to calibration overhead, a similar analysis gives a lost array time of $\tau_{cal} \approx 2T_{cal}(N_t N_m + N_{extra})/(N_t N_m)$ where N_{extra} are the extra antennas taken from the unmoved portion of the array during calibration. The total lost time due to reconfiguration, including both antenna motion and calibration $\tau_{lost} = \tau_{move} + \tau_{cal}$.

Fig 8 shows τ_{lost} as a function of N_m assuming $N_t = 3$; further assumed is that the antenna move time $T_{mov} = 1$ hour (Radford 1999), $T_{cal} = 2$ hours and $N_{extra} = 2$. The lost time is calculated exactly, not using the approximations given above. The dotted line gives the result assuming the move model of Guilloteau (1999) and Yun(1999), the solid line the result with the more realistic move model. This latter curve is almost flat as function of N_m, compared to previous results where it increases with N_m. This is a consequence of two effects, first a more realistic move model which decreases the impact of antenna motion compared to antenna calibration, and secondly a reduced time assumed to move an antenna (1 hour instead of 2). Both effects cause antenna calibration to dominate for the small number of antennas being moved, giving almost constant τ_{lost} for $N_m < 5$.

References

Conway, J.E., 1998, ALMA Memo 216.
Guilloteau, S., ALMA Memo 274.
Holdaway, S.M. Foster, Emerson, Cheng, Schwab, 1996 , ALMA Memo 159.
Holdaway, S.M., 1998, ALMA Memo 1999.
Kogan, L., 1998a, ALMA Memo 202.
Kogan, L., 1998b, ALMA Memo 226.
Radford, S.J.E., 1999, ALMA Memo 280.
Yun, M.S., and Kogan., L., 1999, ALMA Memo 265
Yun, M.S., 1999, ALMA Memo 277