Long baseline analysis of survey data (Neal, June 2015)

1) Fringe rate and delay mapping

I wrote a python script to take some data in AIPS, and produce a fringe rate and delay map. This is basically an FFT of some data as a function of time and frequency, but the tricky bit is getting the alignment on the sky done properly. (If anyone wants the script, let me know).

To check this I simulated some data in AIPS with a 5x3 grid of sources and recovered the map using FRD mapping. The following picture shows one of the baselines to TS001:

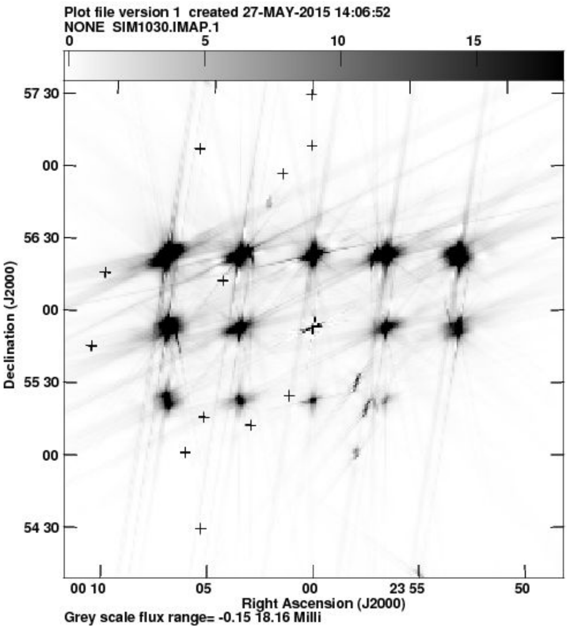

and the next shows the same baseline, but 1

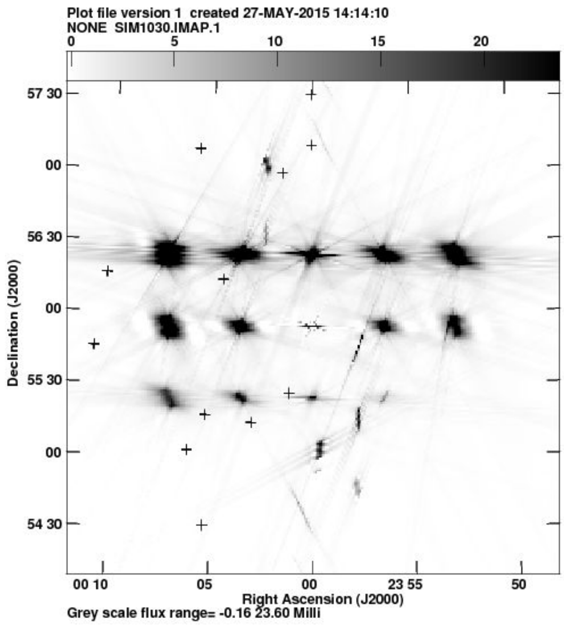

hour later when the uvw's have changed:

So it looks as though the geometry is OK (and the other baselines also seem to work). This has also been done with some of the LOBOS fields, most of which look OK. Here's a typical (OK, one of the best) examples - the raw data, followed by the FRD plot. The crosses are the Wenss sources, one of which seems to be identified as the extra source causing the

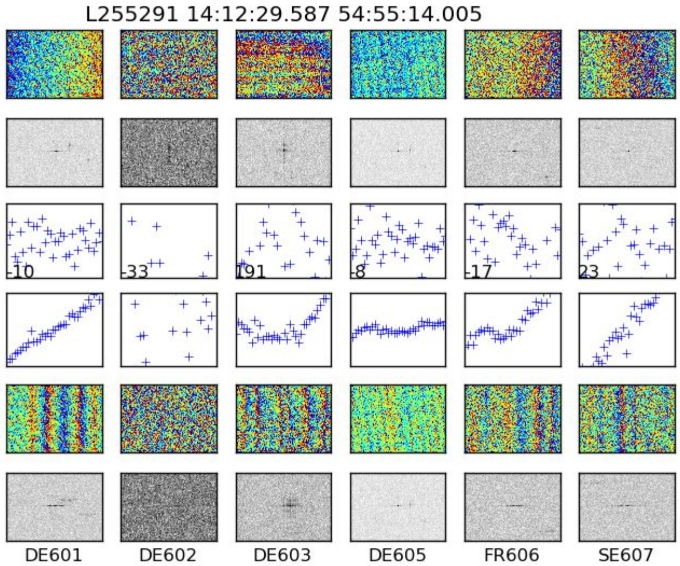

beating.

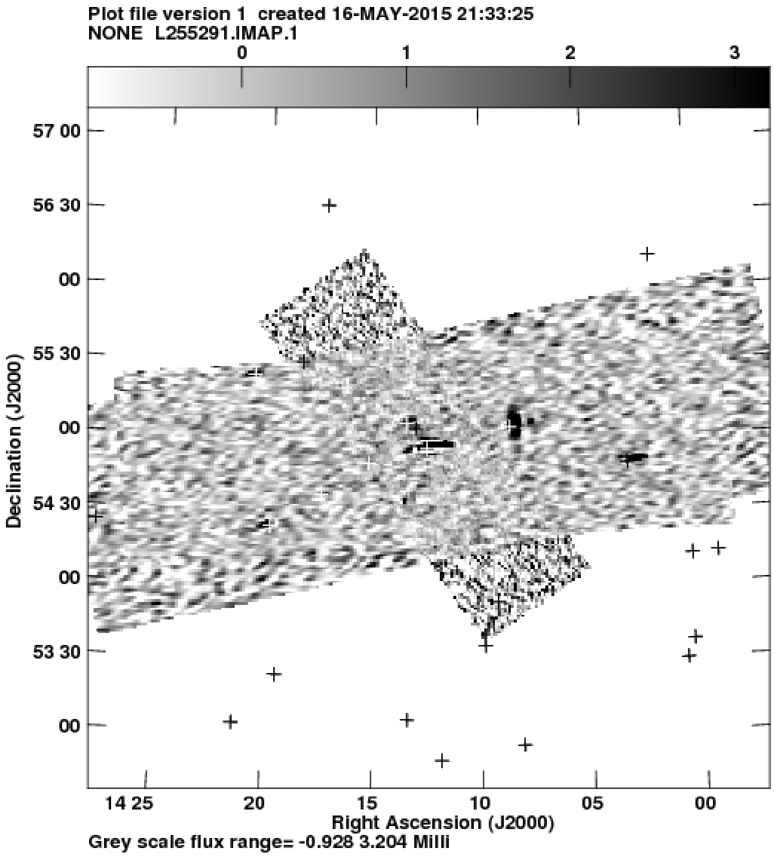

2) FRD mapping of survey field

Tim sent data from a survey field which has a bright calibrator (B1150+497) in the middle of it. It also has 3C266 about 1.4 degrees West of the centre. Due to disk space/time limitations I've only processed the first 50 SB; this was done using locapi (which uses BBS without the beam, and then mscorpol which does have the beam). 3C196 was used to phase up the superterp. The pointing centre was 11:55:41, 49:44:52, i.e. slightly East of B1150+497, but I have shifted the data in locapi to B1150+497. (Opinion: in future I think what you do is not shift the data in NDPPP as this is very time-consuming in terms of multiple data assembly; I think you produce one dataset at the pointing centre and then produce multiple datasets by shifting in AIPS. Revised opinion: this may not work (see below).). The frequency structure of the AIPS data (unaveraged) is 5 IFs, each of 80 24.4-kHz channels (total bandwidth 10MHz). This allows averaging within each IF to shrink the field of view, while still having enough channels (8-10) in each IF to get delay solutions. Here are the fringe rate/delay maps, all on the superterp baselines:

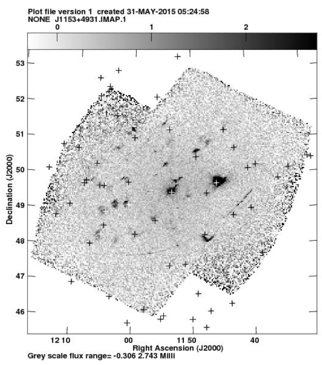 Everything combined	DE601 Effe 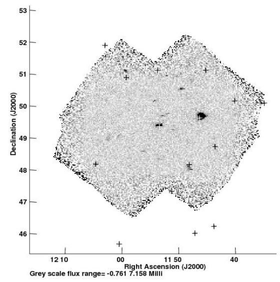 Isberg	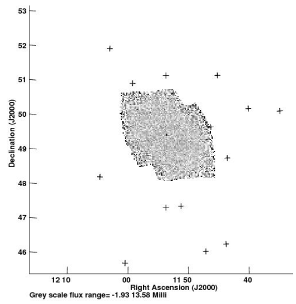 DE602 Unterweilenbach
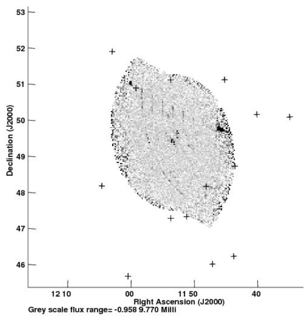 DE603	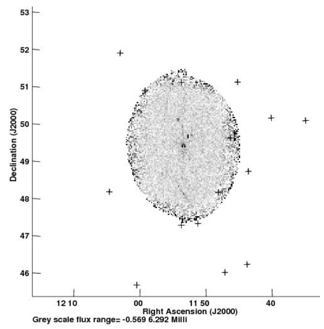 DE604	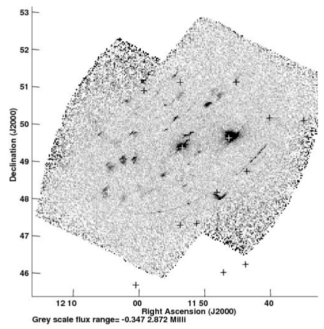 DE605 Juelich
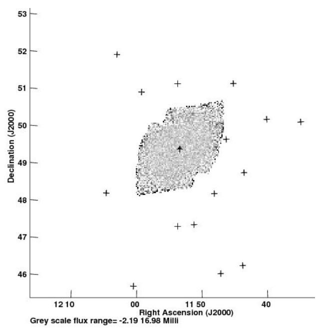 FR606 Nancay	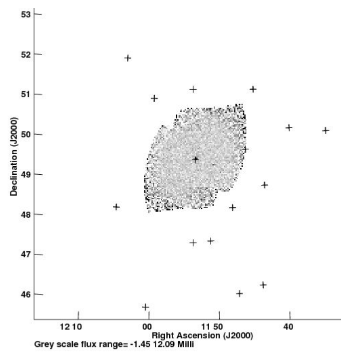 SE607 Onsala	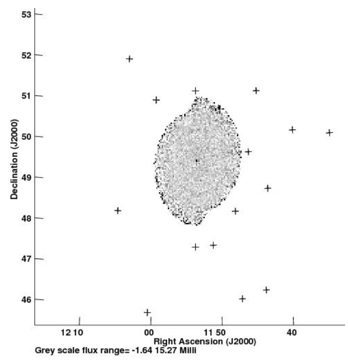 UK608 Chilbolton

Bottom line: you can see B1150+497 on all superterp baselines, but you can only see 3C266 on the German baselines - on longer baselines it is excluded by the time/frequency resolution for a single pointing. Opinion: Effelsberg, Juelich and Norderstedt all being present is going to be vital for serious imaging.

The the fun begins:

1) run UVMOD to reduce the amplitudes by a factor of 1E6 (or AIPS can't handle it)
2) run WTMOD to increase weights by a factor of 1E12 (ditto)
3) produce any shifted datasets (in this case two, at B1150+497 and 3C266). For each of these:
4) SPLAT the data, decreasing the channelisation so that other bright sources are rejected, but so that you still have enough channels for a fringe fit. In practice averaging by a factor of 4 gives 2 smearing lengths which should be safe.
5) FRING using soltyp'L1', zeroing the rates and just keeping delays/phases, using only the IBs. If you try and FRING the RS's you see evidence for strange things happening due to other sources in the field and structure within the sources. With the IBs there is a problem, because both B1150+497 and 3C266 are double sources (cf. below)
6) try to use CALIB for a first go at the amplitudes
7) TACOP the CL table back to the original unaveraged data from 3), then average the data all the way down to 1 channel/IF and 20sec, and try to map it.

5) and 6) will be inaccurate for two reasons: (i) ambiguities in the delays; (ii) B1150+497 has structure. Opinion: I am worried about this, as it implies that for most fields you are going to have to have a model for something in the field, and then refine it in Difmap. I don't even think you can calculate the delays without a good model (see the plots in the next section) as the phase variations resulting from the structures will give artificial ringing in the delay signals.

Squashing the data further in frequency and mapping B1150+497 in Difmap gives an image - cf. the old Merlin image on the right. The spatial frequency coverage is not great in the LOFAR image, and the amplitude calibration is rather iffy, but you see beating due to different components in the German-TS001 baselines. Difmap needs quite a lot of prodding in the right direction (you have to start off with a startmod and use your "imagination")

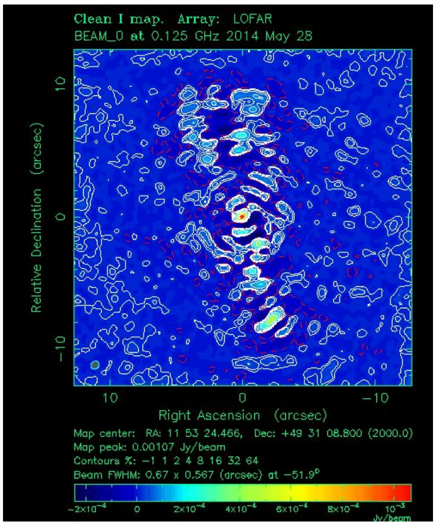

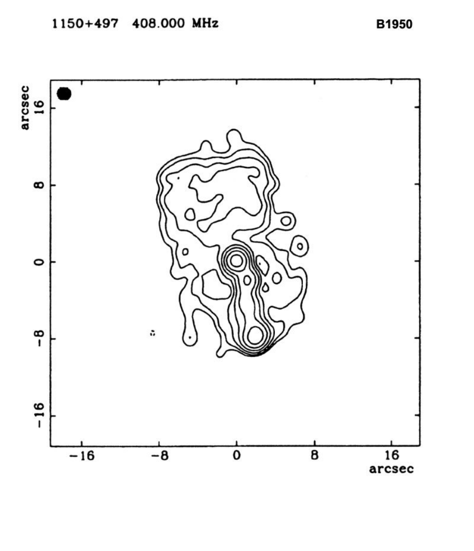

3) Some more detail on the survey field

We have a good test case in these data, because B1150+497 and 3C266 both have simple structure which should be easy to map, and they are much brighter than the rest of the field. So appropriate shifting of the data should give visibility amplitudes which are clearly indicative of the structures, even if the phases need unwinding.

Let's take the Juelich-TS001 baseline. The visibility amplitudes after squashing in frequency to 200kHz (smearing length 0.25 degrees) look like the following:

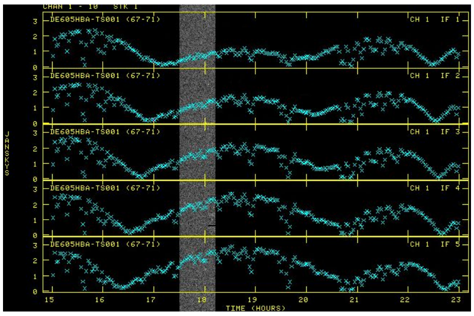

B1150+497

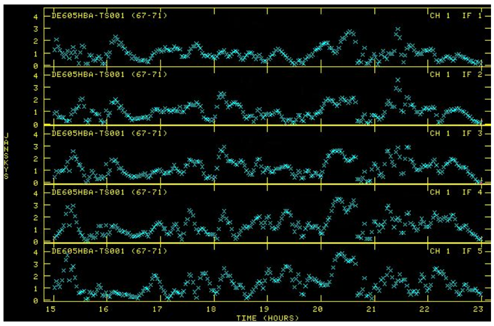

3C266

B1150+497 looks OK - there is a little bit of extra beating which may come out with further averaging, but the structure basically looks well-constrained. 3C266 doesn't - it's a 3" double and should display beating at about twice the frequency of B1150+497 which should be clearly visible. Opinion: this needs to be fixed, as if 3C266 is not easy to do, nothing will be.

The problem with the shifted data on 3C266 is not other sources; if you average it further, there is no improvement in the correlated amplitudes. Specifically, here are the amplitudes and phases for B1150 (left) and 3C266 (right) on TS001 baselines. Each panel is (left to right) Eb,Uw,Ta,Po,Ju,Na,On,Ch, and time increases upwards and frequency to the right in each subplot. (again if you want the parseltongue script, let me know). Thought: it's probably because the beam has been applied in mscorpol, so you can't then shift arbitrarily to other places and expect it to work. Currently producing a dataset at the 3C266 position (awaiting access to a more powerful machine first though). Thought from telecon: probably better if you only use superterp and not all CS for phasing of TS001 - will try this.

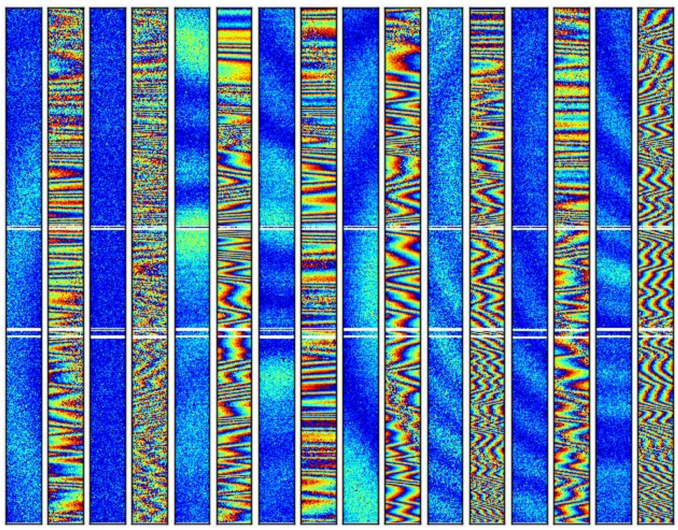

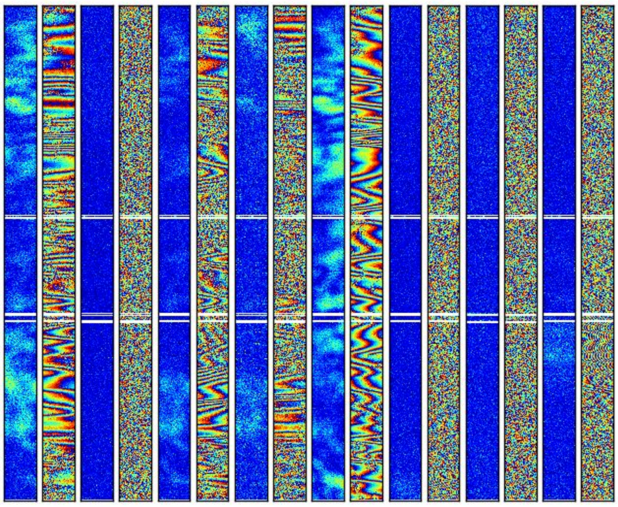

4) Differential smoothing of data

Whatever, it looks like you need to have the ability to stop down the field of view to a smaller area, by smoothing data in frequency and time. This unfortunately needs a smoothing that is dependent on baseline, with shorter baselines being smoothed more to achieve the required level of time and frequency smearing of external sources. The limit on the time smearing is the phase variations of the ionosphere in any direction, and the limit on the frequency smearing is the delays. For ionospheric phase variations on ~1min timescales, and typical delays, you can get a field of view of about 0.5 degree - enough to exclude other sources in the field - only for 50km baselines. (But that's OK).

I wrote a parseltongue script to do this. Here are the results on simulated data, where the simulated field is a N-S double source a few arcseconds apart, together with another bright source 1 degree away. Top is before the adaptive smoothing (to half a degree) and bottom is after it.

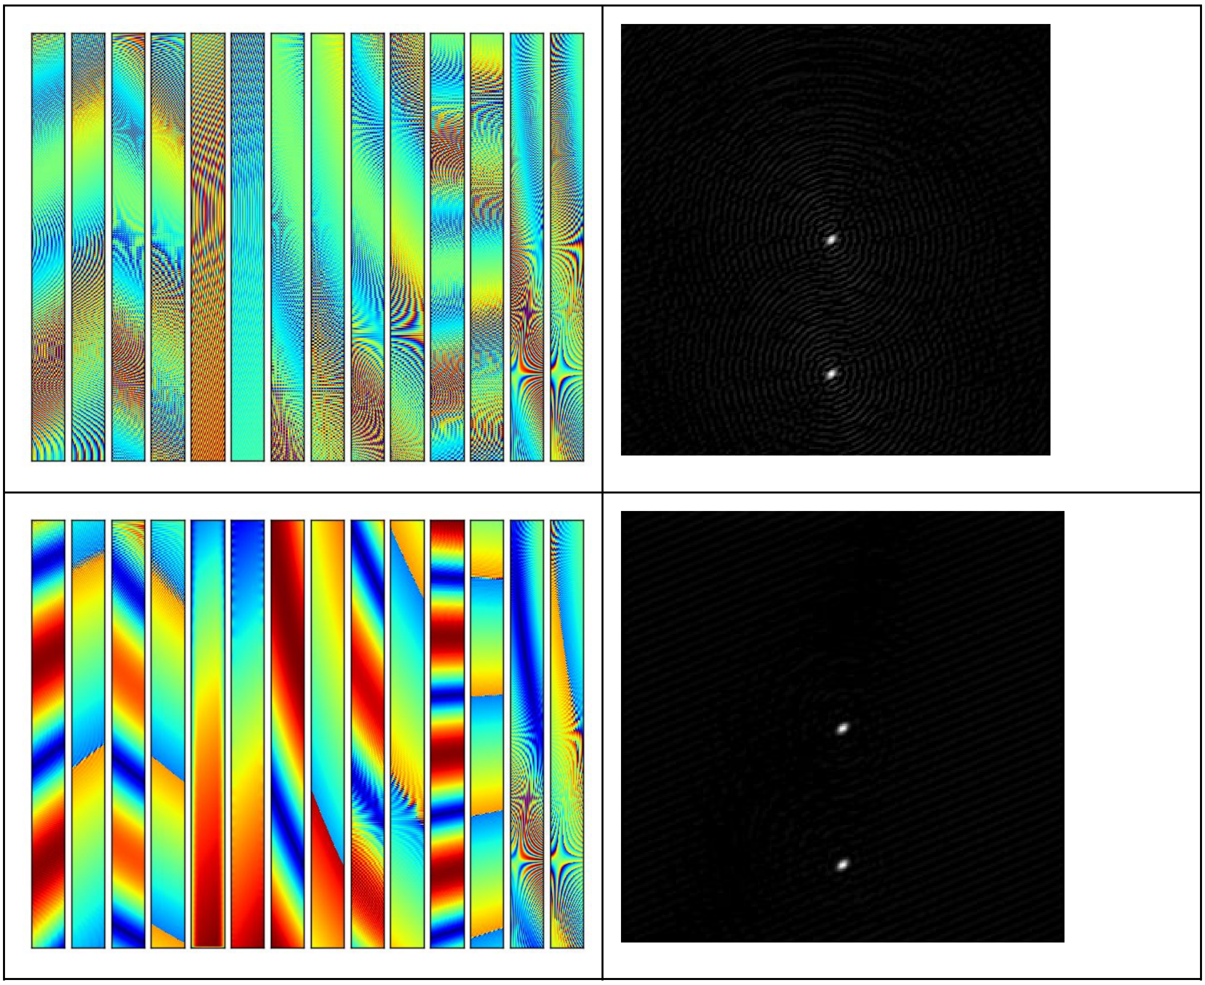

It generally does a good job, except at some points where I suspect the projected baseline is a bit shorter (the program bases the smoothing on the average projected baseline length during an observation). Even on the short Eb-Ju baseline (5th and 6th column) the removal is pretty good. What still needs doing is to flag the band edges where the smoothing gives a falloff in amplitude - this is particularly noticeable once you get down to 50km. It is very, very slow, but I have now managed to parallelize it - this involves multiple threads writing to an AIPS file at once, which is probably a frantically bad idea but seems to work. (Update - Mark Kettenis's view is that this is indeed a frantically bad idea, so the algorithm needs to get cleverer - should we be using hdf5 or something?)

5) New survey field (1340+55)

One subband (200kHz, 8h) of the new survey field (supplied by Tim for the June 2016 workshop) was piped through the fringe-rate and delay mapper with the following result:

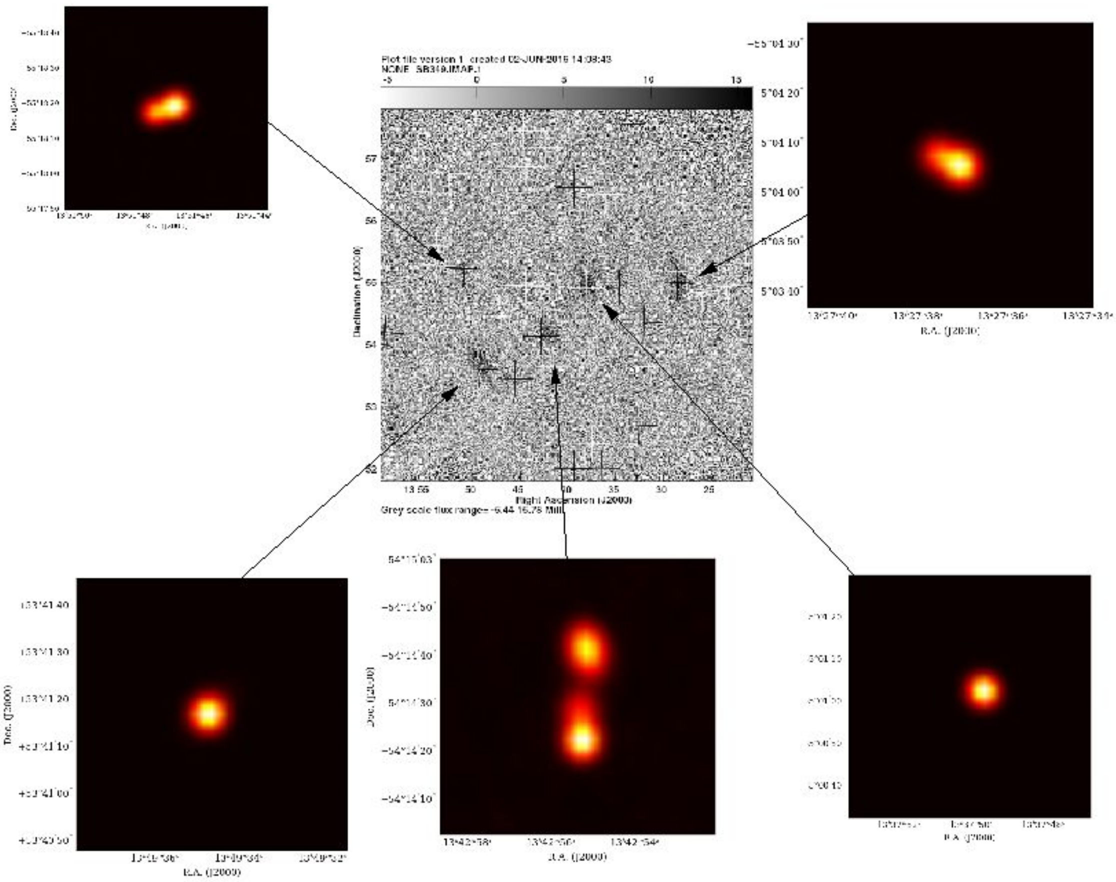

This is the ST001-DE605 baseline. You can see a little on DE601, but not very much - perhaps unsurprising since the sources are ~1Jy rather than ~10Jy. It looks as though there are four or five sources fairly definitely detected in the FRD map (the crosses are WENSS sources, and the colour maps around the edge are from FIRST). All the detected sources are LOBOS “good” calibrators on this baseline. I tried forming ST002 with only the superterp stations, but the results are not as good.

The easternmost source, 1327+5504, is a good LOBOS calibrator (L331908):

L331908 13:27:37.170 55:04:06.200 2015-03-19 00:44:17   PPP-PPPPP

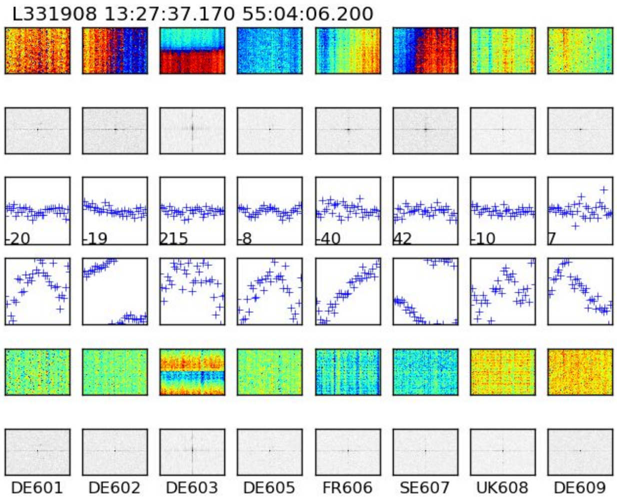

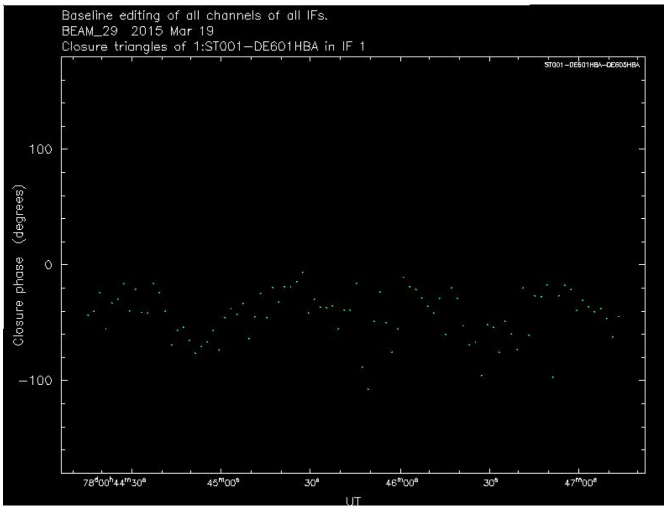

The plot on the left shows the standard LOBOS plot of this source. The phase solutions are all very tight, and the closure phase is coherent. If you take the closure phase over a 200-kHz bandwidth (similar to the survey observations) between TS001, DE601 and DE605 you get the plot on the right: the scatter is bigger because of the smaller bandwidth, but the closure phases look coherent.

To map this source from the survey data, you have to re-form the superterp at the correct position (otherwise it doesn't work). Starting with a point-source model then gives what is obviously the correct map (made by deleting short baselines and using natural weighting):

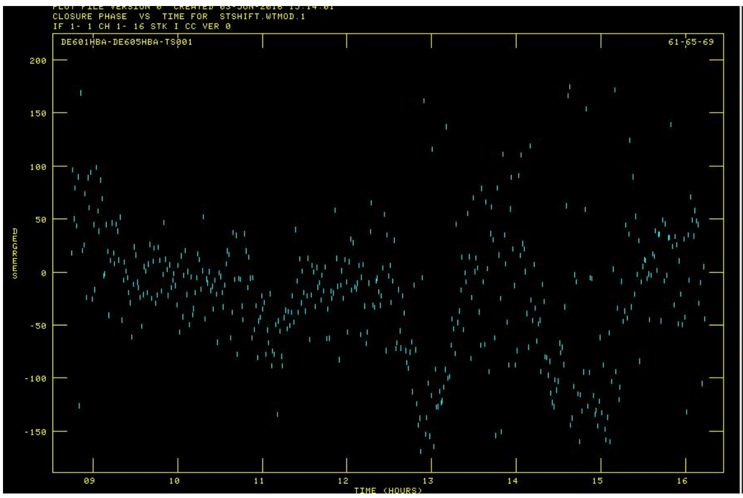

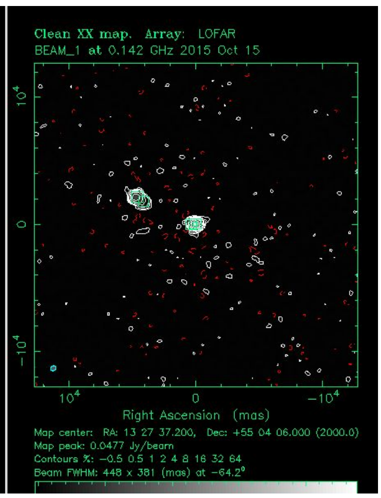

I have written a script to automatically produce small datasets around LOBOS calibrators; it takes (currently) one subband, then forms the superterp at the appropriate place, and averages the data to 20s (probably the maximum safe averaging time, judging by the LOBOS coherence times) and across one subband. It then plots closure phases on a specified baseline - here I've done it for TS001-DE601-DE605. (We don't have Norderstedt I think, which is going to be a pain).

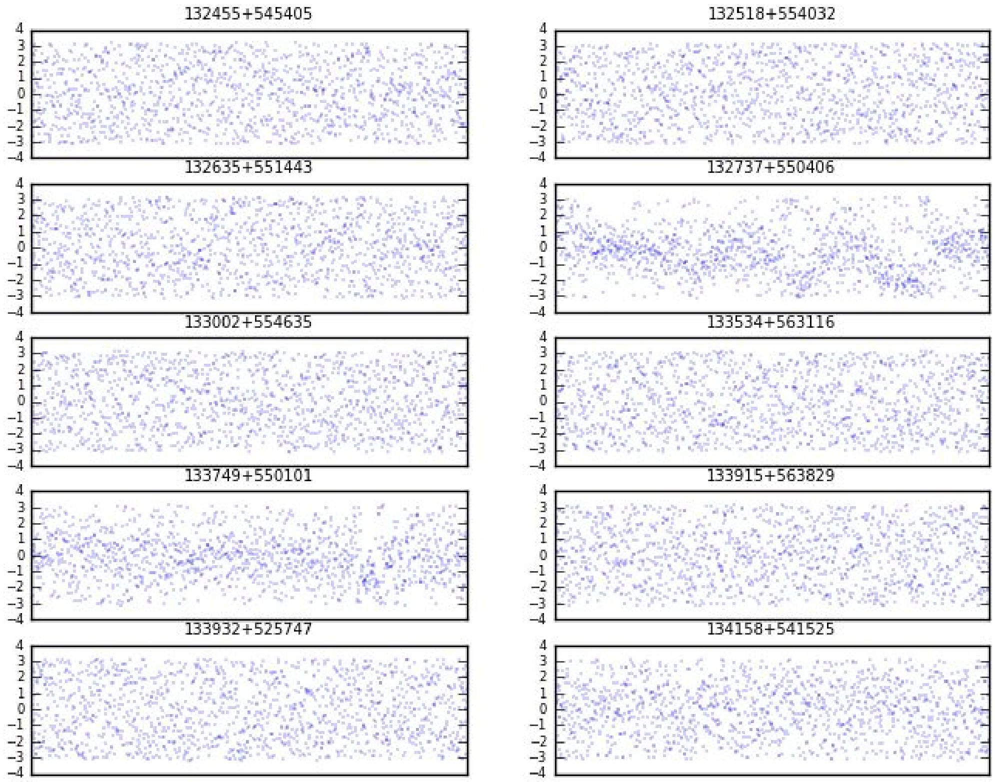

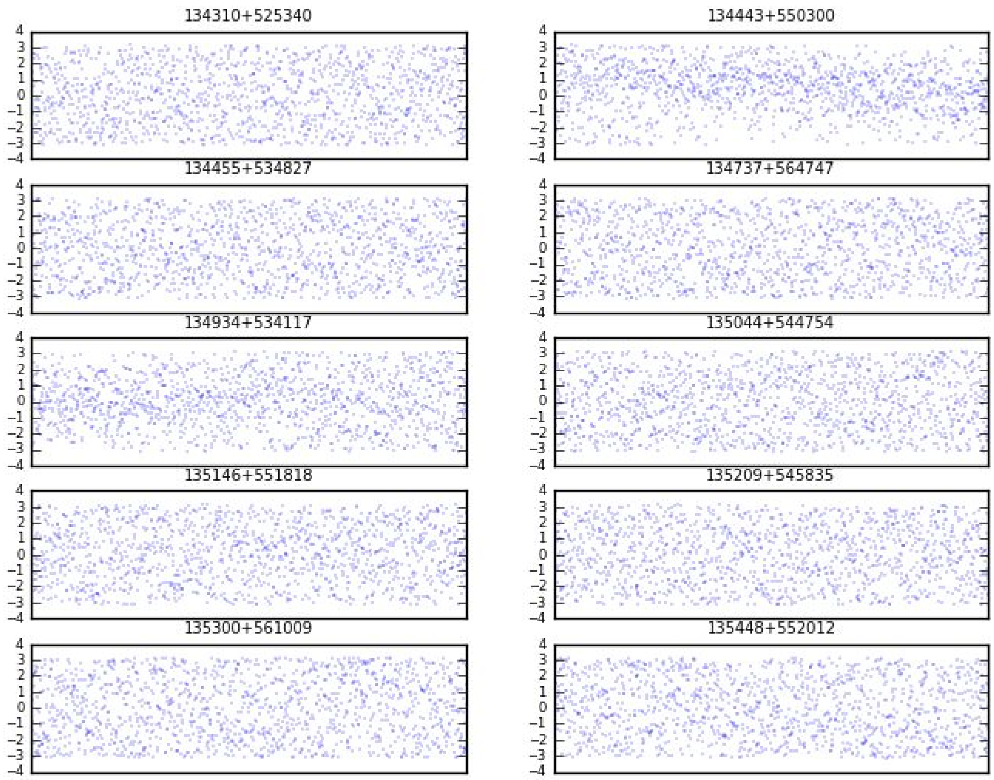

The very good source (1327+5504) stands out here. There are a few other sources whose closure phases are much noisier; if you try to map these, you have to start off with a point source model (as for 1327+5504) but no secondary points appear. So it’s impossible to tell whether you are getting real phase solutions, or just driving the data in the direction of the model. Using these phase calibrators properly is therefore likely to entail using more than one subband, and worrying a bit more about the delays.

Conclusions and opinions (for now)

Just because the source appears on an FRD map doesn't mean you can calibrate on it - calibration depends whether you have good S:N on more than one baseline. You may also need more than 1 subband if it is a LOBOS calibrator, because LOBOS uses bandwidth corresponding to 15 subbands.

I have written scripts so far in a chaotic mixture of NDPPP, casa and aips so far. This needs to be rationalised (probably by somebody who's better at taql than I am).

We need an algorithm for bootstrapping the calibration across the field and from stronger sources to weaker sources, which does not take forever to compute. Maybe it will look something like

a) Make an FRD map of the field on an optimum baseline, probably TS001-DE601. Find the sources which appear. (Easy to automate)
b) For each LOBOS source which can be seen on the FRD map, make an averaged measurement set with one subband and the superterp formed in the correct place. (Automated).
c) For each measurement set where the closure phase on an optimum baseline, probably TS001-DE601-DE605 is clearly visible, make a hybrid-mapping image. (Hard to automate. I have done it by hand in difmap, and you need to be quite careful about not symmetrising the image. It could be done by closure-phase fitting, for which I have a script but which takes a while to run and will only work if there are two components).
d) If no sources are found, repeat b) and c) with increasing number of subbands until a reasonable number of source models are generated, or until delay effects mean that you have to solve for the delays properly. (Painful but in principle easy to automate).
e) At this point, take the best source or sources and solve for the delays using the models generated in parts b)-d). (Easy to automate, but needs AIPS).
f) Begin to use the full dataset at this point. For each LOBOS source which was not previously detected, generate a measurement set with many subbands, apply the delay solution from e), together with the phase solution from the closest detected source, and generate hybrid maps (Slow, and hard to automate. Question: is this the point at which one should stop using the superterp and start using the core stations? Opinion: I don't think so, as the data volume will be prohibitive).
g) For all other sources in the field which are seen in FIRST, use the delay and phase solutions from the nearest source in f) to produce maps. This is a massive undertaking if you have to produce shifted datasets at each point.

Question: do we avoid all this by just throwing all the data at Sagecal?

Preliminary investigations on long baseline BW (done without proper superterp phasing):

If you simply apply the solutions from the bright source 1327+5504, then without further self-calibrations other sources appear using **one subband**:

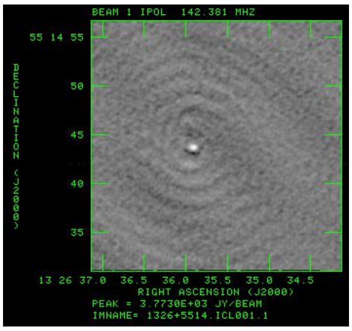

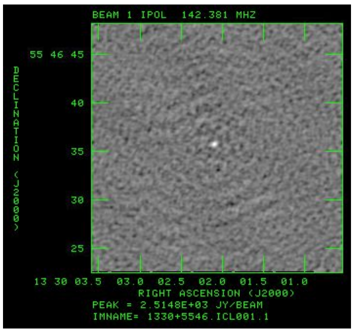

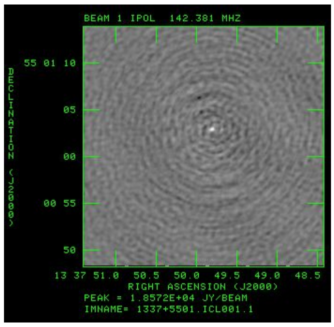

Most of these are relatively nearby (within 1 degree) but the bright source in the middle of the field (1337+5501) also appears. If you tweak the phases up on this you get very marginal coherence on something at the other end of the field, which wasn’t there before

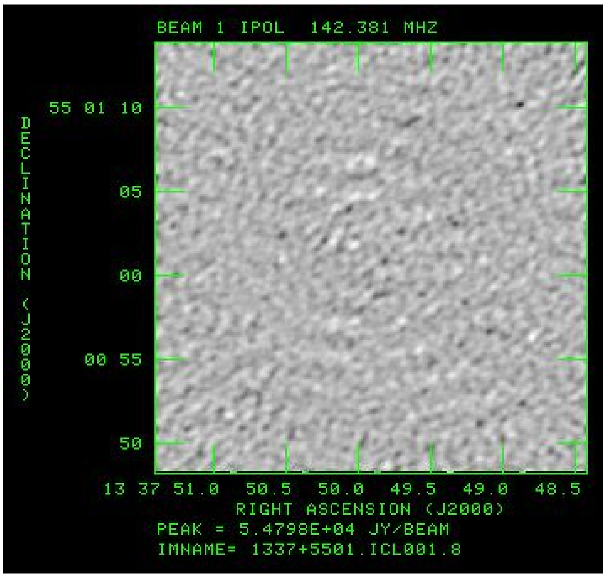

This has also been tried with 12 subbands, and there is now enough bandwidth to do delay solutions:

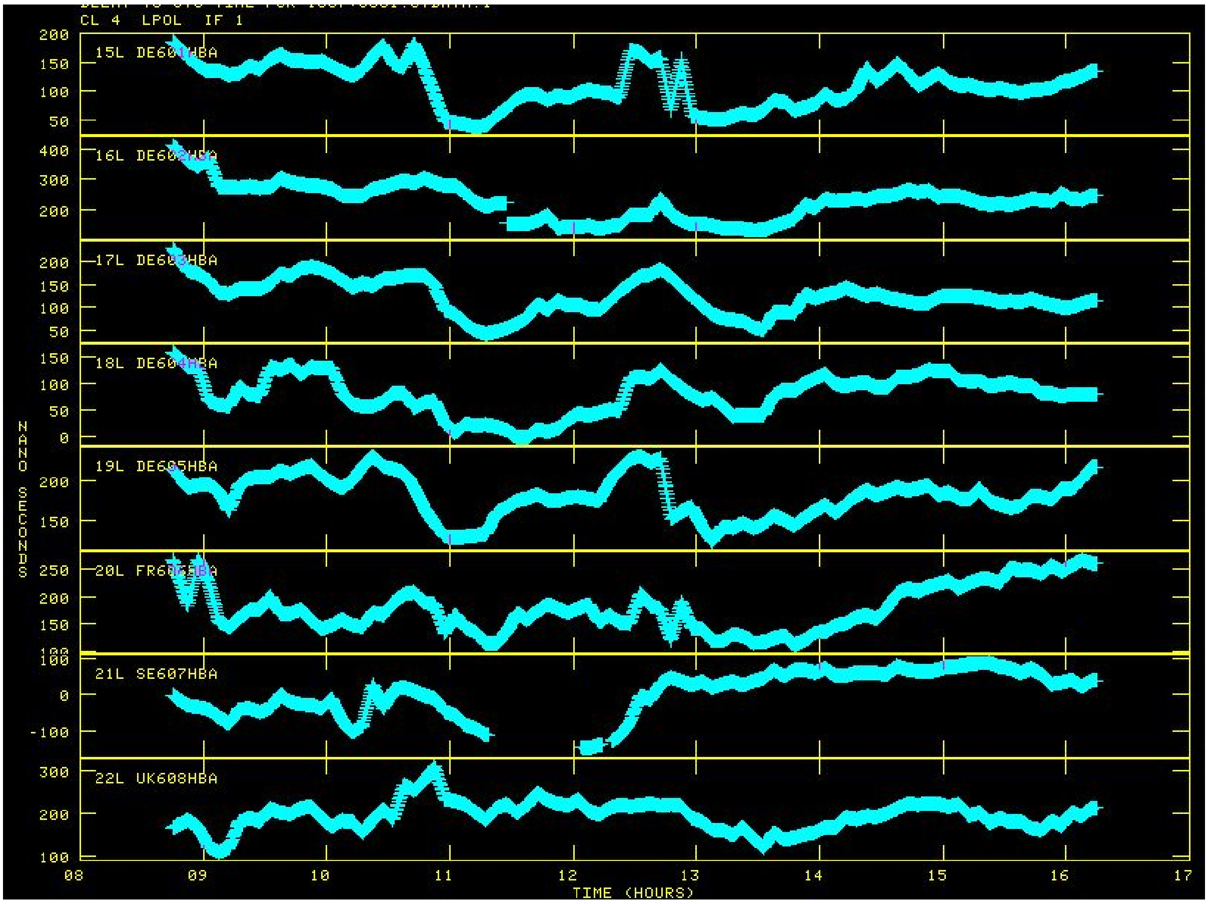

The delay solution is made on the brightest source, using the model together with 5-min solution intervals. Then the phase solution is made, applying these delays, with a 20-sec solution; finally, the amplitude solution with a 5-min solution interval. This can be applied and all the sources are then re-mapped. Six sources are then seen, all of which are in the same part of the field.

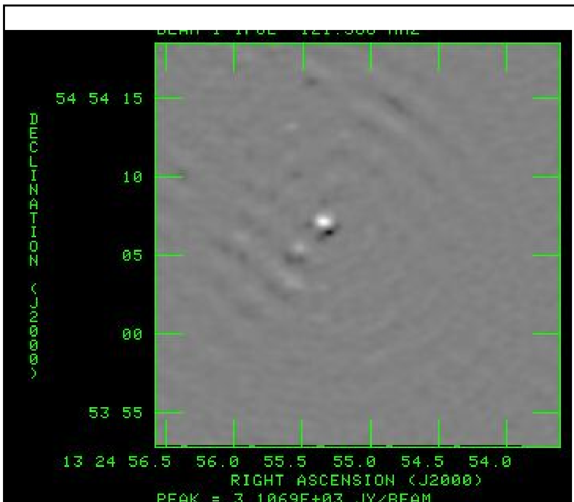

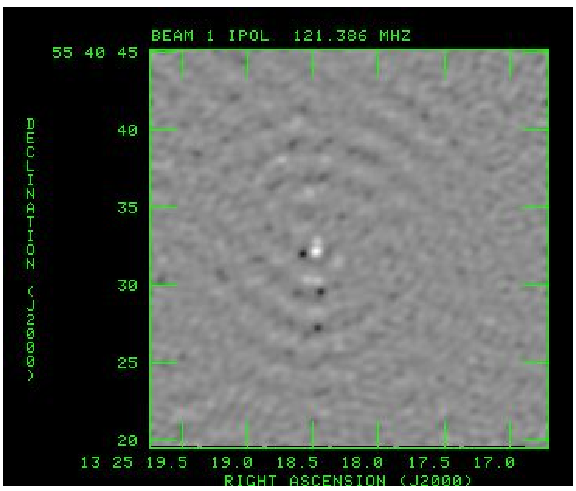

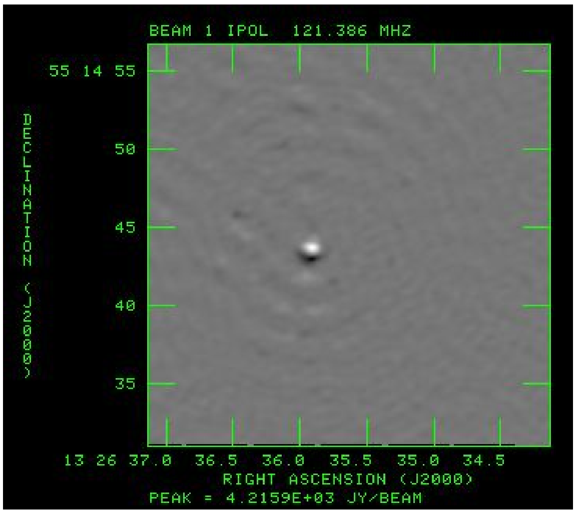

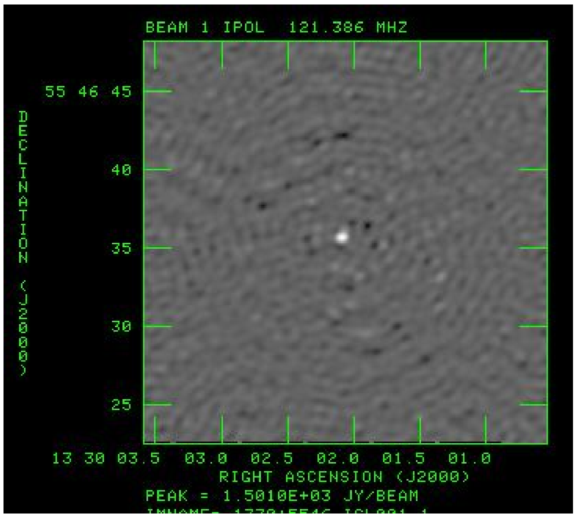

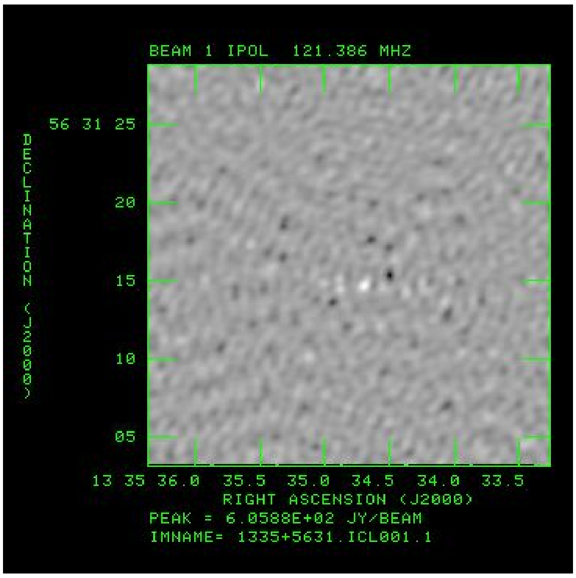

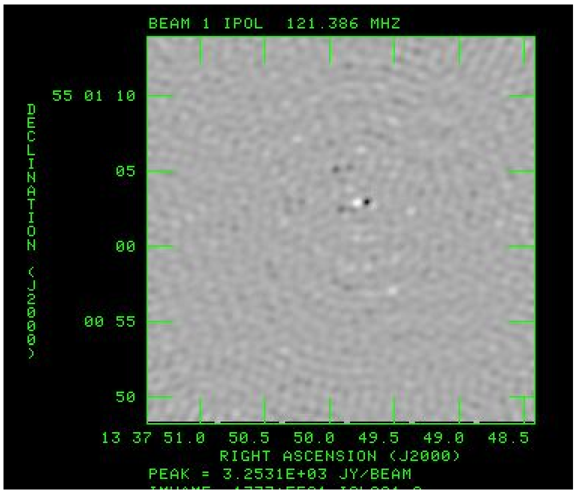

The brightest of these sources, 1324+5454, can then be self-calibrated - phase solutions with 5 minutes look like

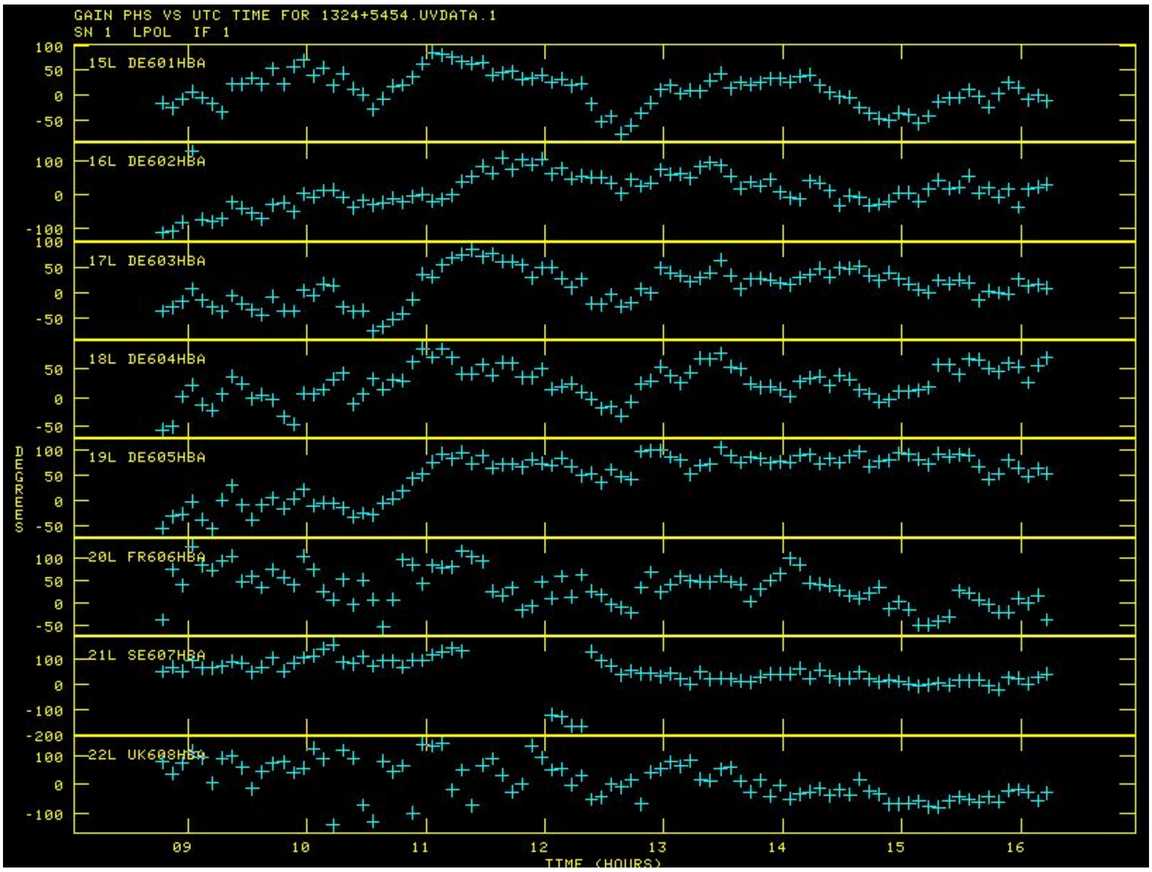

This calibration can then be applied and a new map generated, as follows. I don't like this - it looks like a double source which has been spuriously symmetrised:

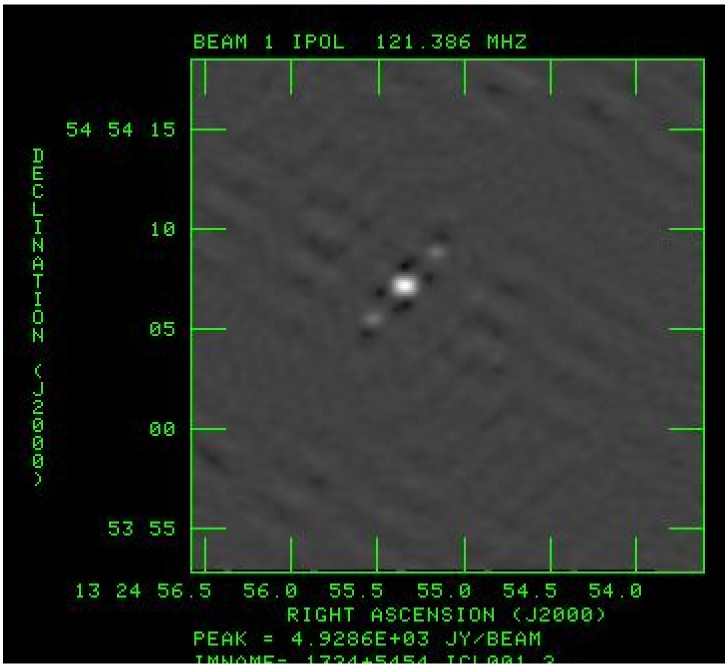

Meanwhile, back at the ranch, we then selfcalibrate the source in the middle of the field (1337+5501) which was marginally coherent in the original calibration on 1327+5504. It now looks a lot better. On the left is the “pipeline” image, and on the right is what I get by hand:

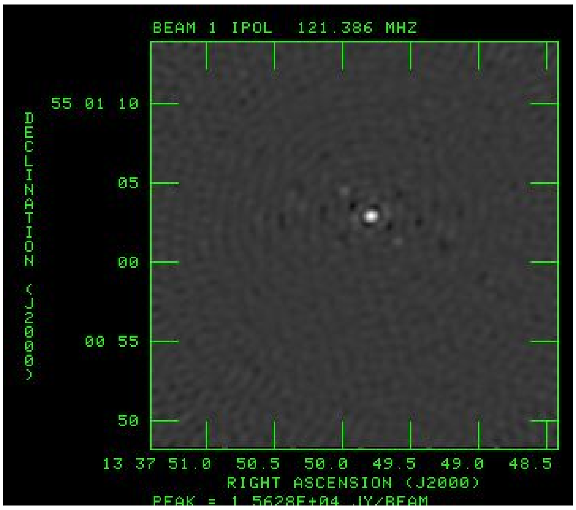

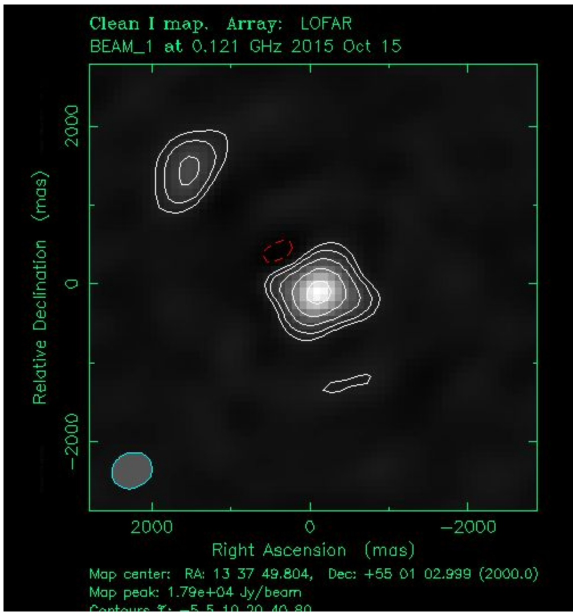

Then we get this calibration and copy it to everything else in the western part of the field that wasn’t imaged successfully before. By this means we get three extra sources in the western part of the field:

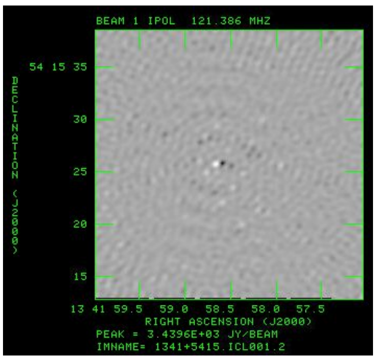

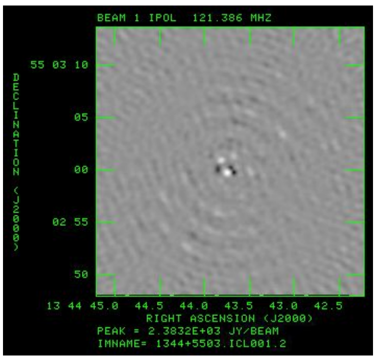

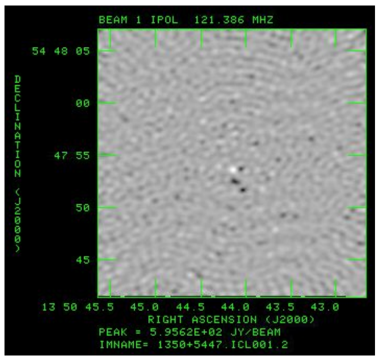

In principle, you can do hybrid mapping using AIPS. However, the central source has been done with love using Difmap, using several cycles of phase self-cal, followed by a gscale and then amplitude selfcal on times from 30 min down to 5 min, using baselines >100klambda, giving the following:

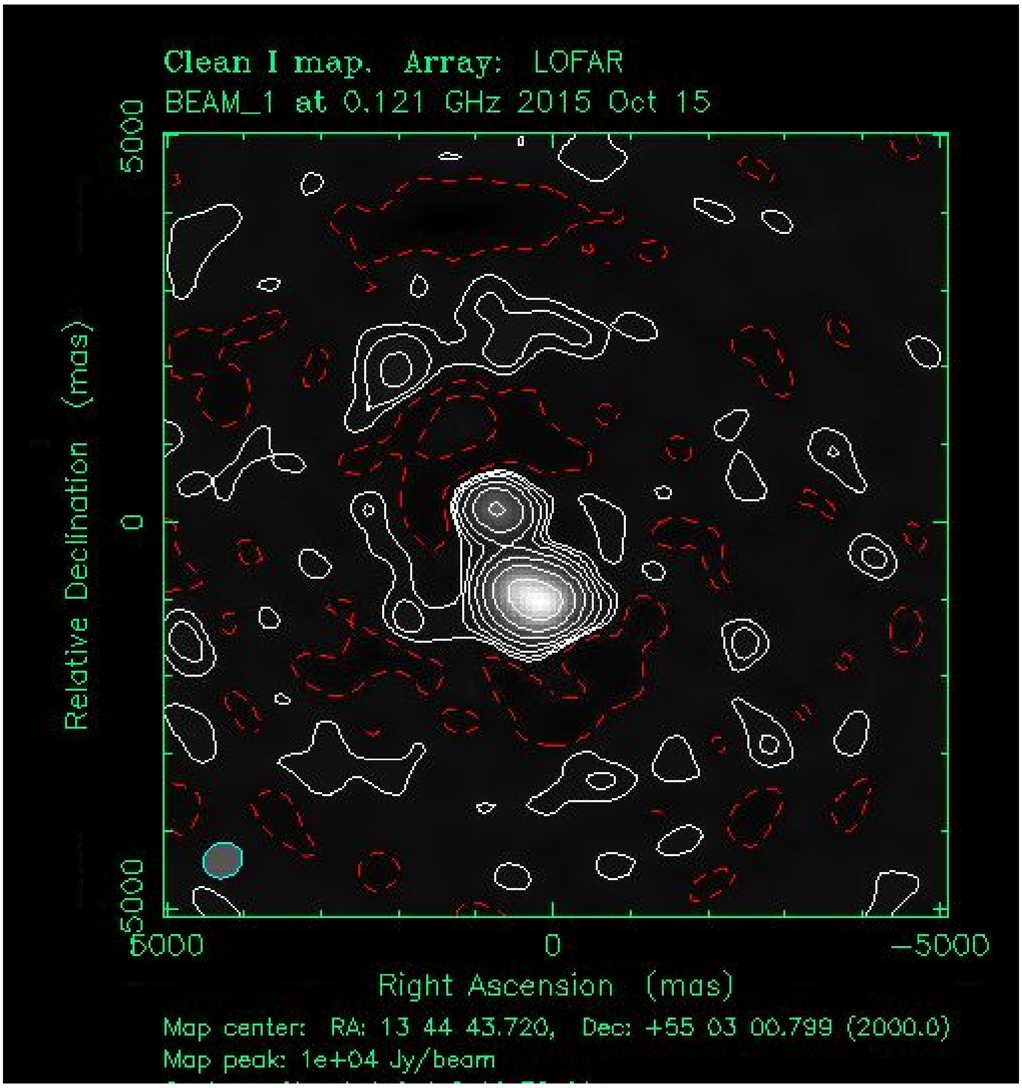

Mapping all 21 sources in the field gives 16 good maps (point appears immediately with coherent phases) and 5 not-so-good (not confident all structure is real);

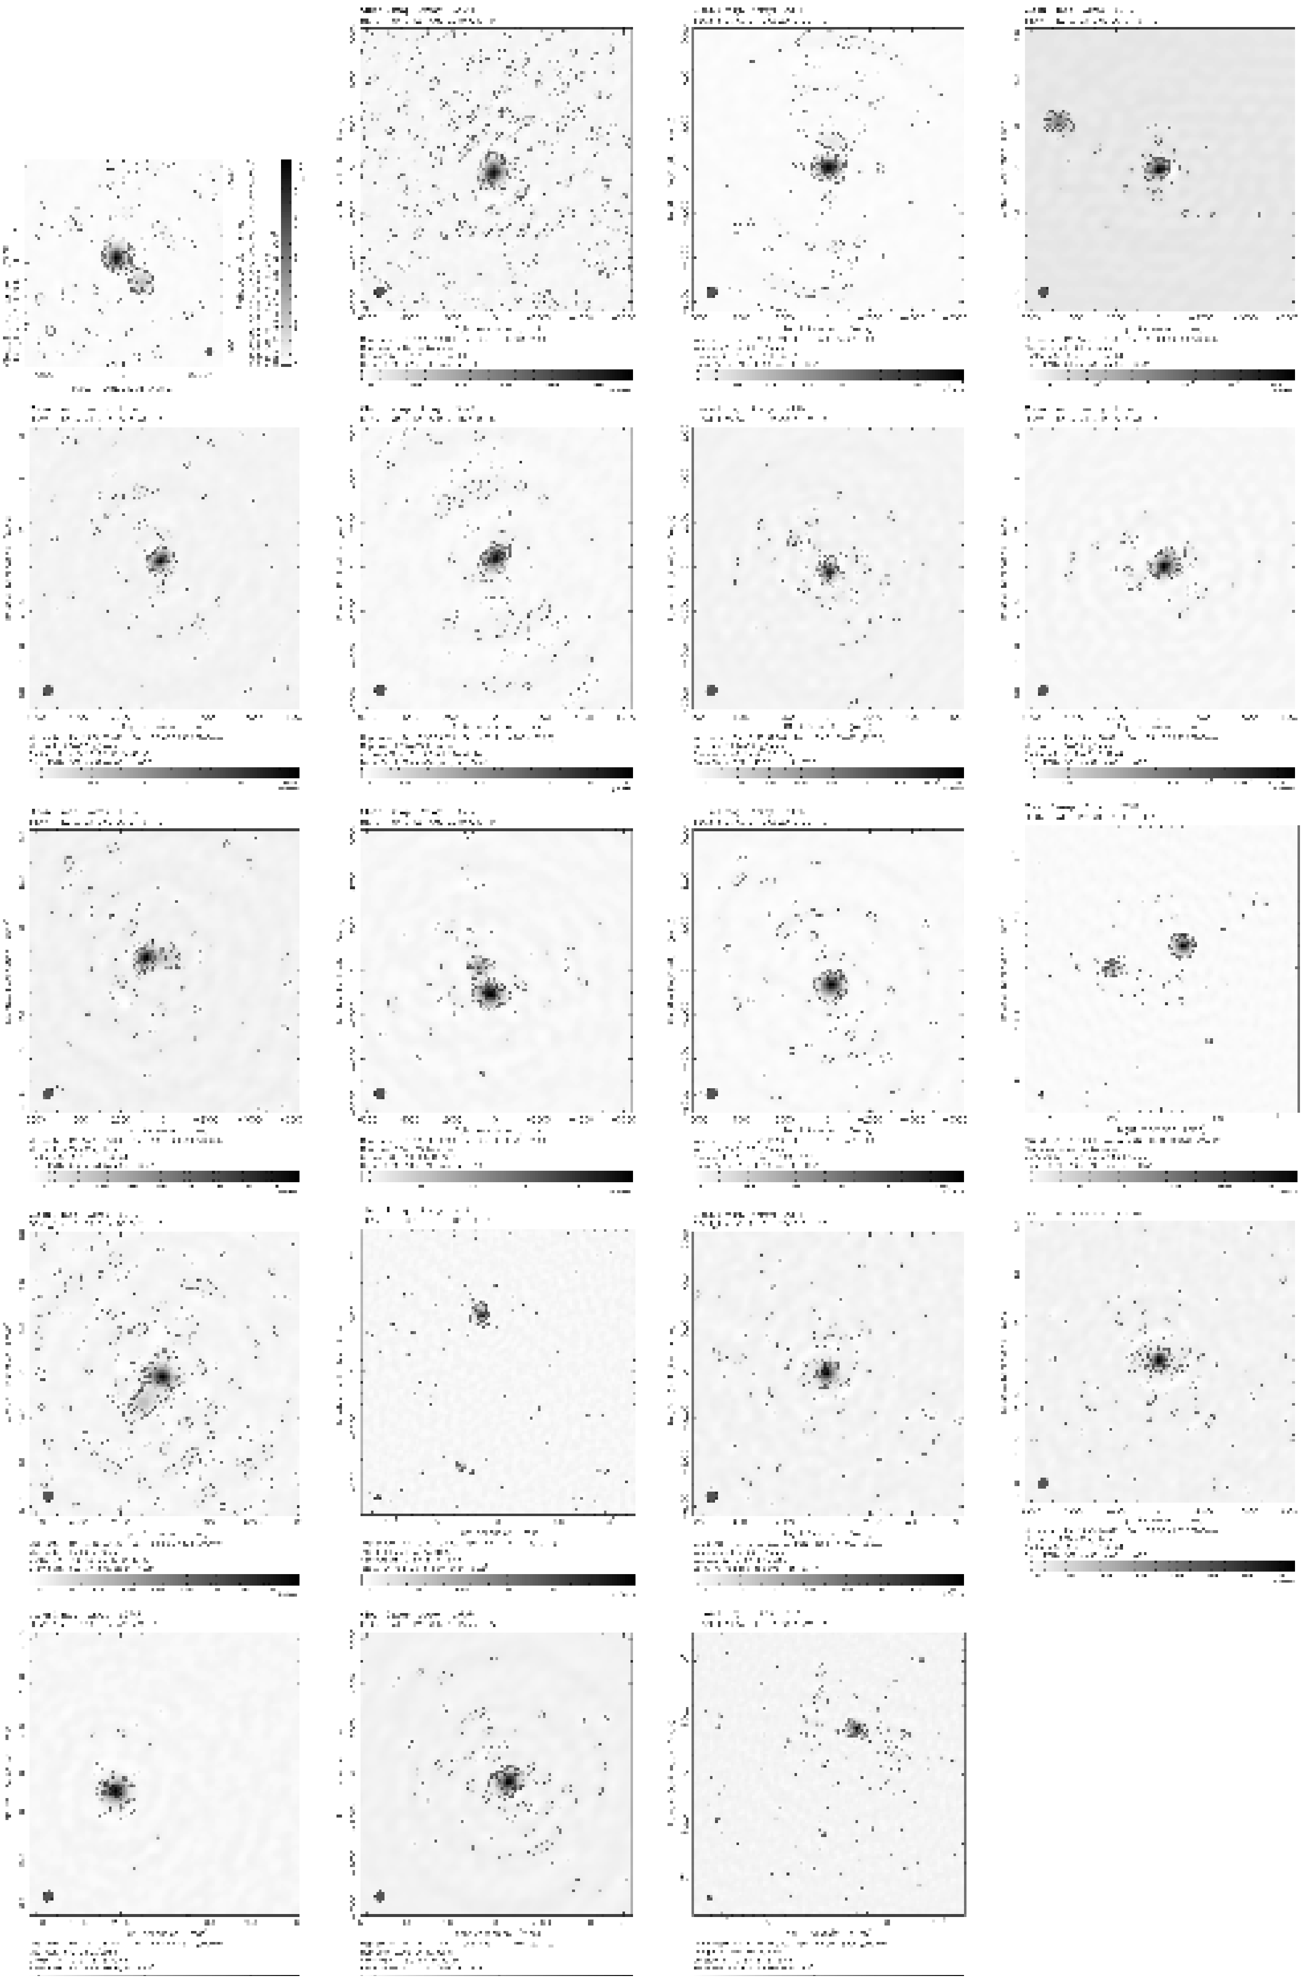

Some conclusions and opinions:

a) In a typical field (assuming this is typical), there will be a source whose closure phases are visible in a single subband. This seems to require a few Jy. Using this, it is possible to get ~10 sources in the field with enough coherence to derive phase solutions.
b) This can be done with about 12 subbands. In a relatively small surveys field, without flanking fields, the brightest source may require an explicit fringe-fit for a delay solution. This can be done as a side-branch if necessary and imported back via a parmdb.
c) It is possible to write this into an algorithm.
d) The delay solutions are going to be needed for wideband imaging, but delay solutions on only one source in the field are probably needed; phase solutions can then be used to move around the field effectively.
e) The tricky bit is to make optimal images of each source. It seems to be a waste not to do this, given the superb quality of the images potentially available. This is harder to automate, but worth thinking about carefully.

6) Latest progress with LOBOS (now >20k sources)