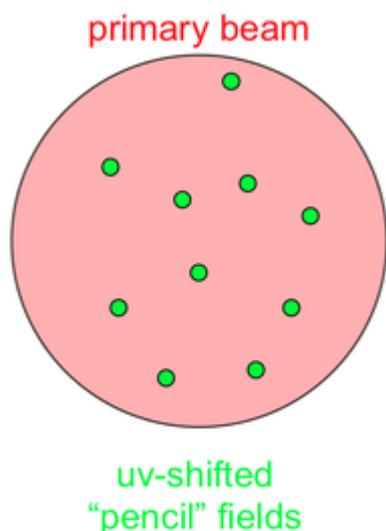


# Multiple phase centres in DiFX2.0

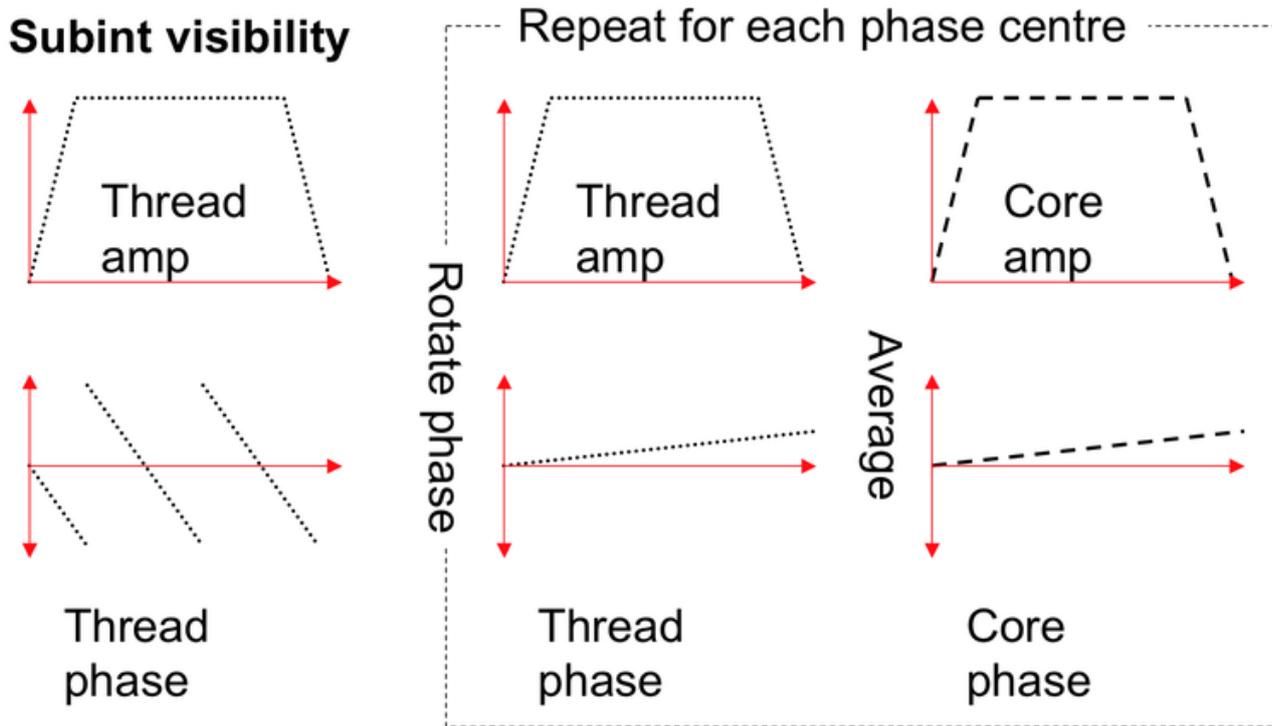
This is the biggest news in DiFX2.0. In a single pass, you can generate a near-arbitrary number  $N$  of independent phase centres, each of which are written out as an independent data file, just as if you had run the correlator  $N$  times.

The image below illustrates what is happening. The red circle indicates the primary beam of the antenna - something like 30 arcmin for a 25m dish at 1400 MHz, for example. However, imaging that entire primary beam at VLBI resolution is prohibitive - it requires 10s or 100s of TB of visibilities to be produced, and the end result would be an image with 100s of Gpixels, 99.9999999% of which would be noise. Since we don't want that, your typical VLBI correlator produces visibilities with time resolution of  $\sim$ seconds, and frequency resolution of  $\sim$ 100s kHz. This means that the visibility data set is much smaller - only GBs. Also, time and frequency smearing limits the useful field of view to a few arcseconds, meaning the final images are much smaller (and much better matched in size to the objects typically studied). A single correlator field of view (a "pencil beam") is represented on the image (not to scale) as a green circle.



Now, there are a few ways to get dozens of these pencil beams. In any old correlator, including any version of DiFX, you could run dozens of correlator passes, each with a different phase centre. That obviously gets prohibitive in time pretty quickly. Alternatively, you could write out the data at high time and frequency resolution (avoiding the smearing effects) and uvshift post-correlation. Not many correlators can produce such fine-grained data - DiFX and the JIVE correlator are the only two I know of. Still, to do shifts to the edge of the primary beam, you need to turn 100s of TB of baseband data into many 10s of TB of visibilities, then repeatedly read that dataset into memory, shift and average, and write it out again. Again, this rapidly becomes prohibitive in time (and disk space!)

If, on the other hand, you hit the shifting and averaging operation *inside* the correlator, you avoid all that nasty IO cost of writing to and reading from disk. So that is what we've done for DiFX2.0. Using [vex2difx](#), you can specify as many additional phase centres to be correlated for a given source as you like. A model will be computed for each of them and passed to the correlator; the correlator periodically differences the model for each desired phase centre and the model for the pointing centre and applies a uvshift. After uvshifting, the need for high spectral resolution goes away, and the data can be heavily averaged in frequency. This is illustrated in the image below.



By default, this differencing, phase rotation and averaging will be carried out once per subintegration. Since subintegrations are typically hundreds of ms long, and are usually divided up between 4 or 8 threads or so, this means the shifting is typically being done every 20ms or so. This is usually good enough for even the most outrageously large uvshifts. However, in case it isn't, the facility exists to do the shift and average more frequently. See the [vex2difx](#) documentation for information on how to force more frequent uvshifts.

So, as you can see, this mode of operation requires much more memory than usual. On one hand, you have the compute threads, which have much higher spectral resolution than usual, but still just for one copy of the visibilities per baseline. Then, back at the main thread, the spectral resolution is normal, but you have as many copies of the visibilities as you have phase centres. So, some care is needed in parameter selection to ensure that you don't try to allocate too much memory. If you start swapping memory, performance will nose-dive.

As shown on the [benchmarks](#) page and explained on the efficiency improvements pages [here](#) and [here](#), various changes to the internals of DiFX2.0 have made this a very efficient process. Doing 100 phase centres only takes 10% more time than a single phase centre at the same high internal spectral resolution (!!). Of course, no sane person would run a continuum observation at such high internal spectral resolution, so comparing to a "normal" continuum observation, 100 phase centres takes 2.3 times longer (or 2.3 times as much computational power).

From: <https://www.atnf.csiro.au/vlbi/dokuwiki/> - **ATNF VLBI Wiki**

Permanent link: <https://www.atnf.csiro.au/vlbi/dokuwiki/doku.php/difx/difx2.0multipc>

Last update: **2015/10/21 10:08**

