

This page is redirected to [benchmarks](#).

DiFX Benchmark Results

It is the intention of this page to provide a means of documenting historical benchmark results. Each different dataset that anyone has used will form a subsection, consisting of a table containing the cluster used, the correlation parameters, the DiFX version and the time taken.

If you are interested in a description of how to run benchmark tests, please see the [benchmarking documentation](#).

TC016A - module-based correlation

This was correlated with the current VLBA software correlator cluster (10 nodes, dual quad core Intel Xeon 5420 @ 2.5 GHz, 6MB shared L2 cache/CPU, 7 compute threads/node), with the data played back off module, allowing read speeds up to ~950 Mbps/station. It was a 9 station experiment, and the primary goal was to compare DiFX1.5 with DiFX2.0 for large numbers of channels, and to benchmark the multiple phase centre code in DiFX2.0. The dataset was 60 seconds long, and all were correlated in full polar mode. Note that for DiFX1.5.2 the “true” time should really be boosted by ~4% because the correlator shuts down a little early and ditches part of the last integration.

DiFX version	# spectral points	# FFTs buffered	# phase centres	Time (s)	Notes
DiFX-1.5.2	128	1 (N/A)	1	60	
DiFX-1.5.2	256	1 (N/A)	1	63	
DiFX-1.5.2	1024	1 (N/A)	1	222	
DiFX-1.5.2	4096	1 (N/A)	1	295	
DiFX-2.0	16	1	1	56	
DiFX-2.0	128	1	1	54	
DiFX-2.0	256	1	1	54	
DiFX-2.0	1024	1	1	212	
DiFX-2.0	4096	1	1	290	
DiFX-2.0	16	10	1	54	
DiFX-2.0	128	10	1	54	
DiFX-2.0	256	10	1	66	
DiFX-2.0	1024	10	1	83	
DiFX-2.0	4096	10	1	157	
DiFX-2.0	16	25	1	67	
DiFX-2.0	4096	25	1	150	
DiFX-2.0	4096	10	10	154	
DiFX-2.0	4096	10	30	167	
DiFX-2.0	4096	10	100	178	

DiFX-2.0	4096	10	500	405	Mostly due to disk write speed limitations. "Correlate" time was 244 seconds
----------	------	----	-----	-----	--

DiFX2.0 is slightly faster for "normal" correlation, but much faster for large numbers of channels when FFT buffering is turned on. The scaling with number of phase centres is that 500 additional phase centres less than doubles the time taken for one phase centre, at least when comparing like numbers of spectral channels and neglecting the impact of writing to disk (which since these tests has been made much more efficient and should now no longer be a factor). Thus, the main effect that needs to be taken into account when processing multiple phase centres is simply the cost of going to higher spectral resolution, which is a cost of ~x3.

TC016A - file-based correlation

This dataset is not yet available from the ftp area. It consists of 8 VLBA stations, 512 Mbps data (8x16 MHz bands). Some ~15 minutes of data is available in total - for the tests described below, a subset consisting of 60 seconds was used.

Cluster	# compute nodes	# threads/node	DiFX version	# spectral points	# phase centres	Time (s)	Notes
VLBA	10	2	1.5.1	128	1	160	Full polar
VLBA	10	2	2.0	128	1	146	Full polar, no FFT "batching"
VLBA	10	2	1.5.1	1024	1	240	Full polar
VLBA	10	2	2.0	1024	1	202	Full polar, no FFT "batching"
VLBA	10	2	2.0	1024	1	175	Full polar, 10 batched FFTs
VLBA	10	2	1.5.1	4096	1	400	Full polar
VLBA	10	2	2.0	4096	1	402	Full polar, no FFT "batching"
VLBA	10	2	2.0	4096	1	300	Full polar, 10 batched FFTs
VLBA	10	2	2.0	4096	100	327	Full polar, 10 batched FFTs
VLBA	5	4	1.5.1	128	1	164	Full polar
VLBA	5	4	2.0	128	1	145	Full polar, no FFT "batching"
VLBA	5	4	2.0	128	1	141	Full polar, 10 batched FFTs
VLBA	5	4	1.5.1	1024	1	317	Full polar
VLBA	5	4	2.0	1024	1	262	Full polar, no FFT "batching"
VLBA	5	4	2.0	1024	1	184	Full polar, 10 batched FFTs
VLBA	5	4	1.5.1	4096	1	570	Full polar

VLBA	5	4	2.0	4096	1	598	Full polar, no FFT "batching"
VLBA	5	4	2.0	4096	1	384	Full polar, 10 batched FFTs
VLBA	5	4	2.0	4096	1	375	Full polar, 25 batched FFTs
VLBA	3	7	1.5.1	128	1		Full polar
VLBA	3	7	2.0	128	1		Full polar, no FFT "batching"
VLBA	3	7	1.5.1	1024	1		Full polar
VLBA	3	7	2.0	1024	1	534	Full polar, no FFT "batching"
VLBA	3	7	2.0	1024	1		Full polar, 10 batched FFTs
VLBA	3	7	1.5.1	4096	1		Full polar
VLBA	3	7	2.0	4096	1		Full polar, no FFT "batching"
VLBA	3	7	2.0	4096	1		Full polar, 10 batched FFTs

It should be noted that the 1.5.1 tests with the VLBA required some estimation, since 1.5.1 could not skip past the start of disk files and took ages to get to the right spot in the files, messing up the timing. Note that the advantage of DiFX2.0 in large numbers of channels becomes more pronounced when the nodes have more threads (which reduces the amount of L2 cache each thread has). For the same reasons, the advantage of DiFX2.0 would be greater for experiments with more antennas (there are usually 10 VLBA antennas, not 8). Basically, DiFX2.0 holds the scaling apparent here much better to larger numbers of channels, Cores, or antennas, much better than DiFX1.5.1 can. It should also be noted that the VLBA cluster only has 512 MB of RAM per CPU core. For large numbers of channels, this is getting uncomfortably tight. For clusters where high spectral resolution is expected to be commonplace, 1GB of RAM per core might be better. This is probably standard by now anyway.

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

Permanent link: <https://www.atnf.csiro.au/vlbi/dokuwiki/doku.php/difx/benchmarks?rev=1323356539>

Last update: **2011/12/09 02:02**

