

DFT and FFT implementation in DiFX

Previously mpifxcorr utilised FFT(Fast Fourier Transform) functions from Intel's IPP(Intel Performance Primitives) libraries however this restricted the user to setting the channels to a power of two. However the dft functions (which allow arbitrary-length FFTs) are highly optimised (see "[Short benchmark of Multi-threaded FFTW vs Intel IPP](#)" for a comparison with IPP FFTs and fftw). For the specific case of DiFX there is little difference in the time taken to perform a DFT compared with an FFT, particularly for powers of 10 and presumably for powers of 3, 5 etc (see below). The trunk version of mpifxcorr now supports the calling of the dft functions when nFFTChan is set to a number which is not a power of two. The behaviour of DiFX should be entirely unchanged from previous versions if nFFTChan is set to a power of 2 (the fft routines are used as before).

Setup

The default values of stride length (strideLength) and xmac stride length (xmacLength) are based on the assumption that the number of channels (nFFTChan) is a power of 2. At the moment there is no automated feature to set these values based on the number of channels. Thus they have to be set manually in the SETUP default section of the v2d file. However they must all satisfy the constraint that they are factors of the nFFTChan. Setting all four of the variables below is advised for now.

eg

```

SETUP default
{
    nFFTChan=503
    nChan=503
    strideLength=503
    xmacLength=503
}
...

```

Performance

[NB DiFX-2.0 terminology used in this section]. For performance a number of tests were performed using vlbi_fake to simulate 6 data streams. These tests used nChan's of 16, 500, 503, 512, 1000, 1024. Since vlbi_fake sends data limited to the speed mpifxcorr can process it, thus can be used to benchmark mpifxcorr's performance for different nChans and specAvg. The data can be organised into three sets, first where there are no real comparable features, next there was no averaging (specAvg = 1), and finally where specAvg was set such that the number of output channels was comparable.

Set 1

nChan	specAvg	nChanOut	Mean data transfer rate
16	16	1	442.39 Mb/s

nChan	specAvg	nChanOut	Mean data transfer rate
500	10	50	278.52 Mb/s
512	16	32	277.71 Mb/s
1000	20	20	279.19 Mb/s
1024	16	64	237.46 Mb/s

Set 2

nChan	specAvg	nChanOut	Mean data transfer rate
16	1	16	438.25 Mb/s
500	1	500	277.84 Mb/s
503	1	503	214.82 Mb/s
512	1	512	276.31 Mb/s
1000	1	1000	234.61 Mb/s
1024	1	1024	235.05 Mb/s

Set 3

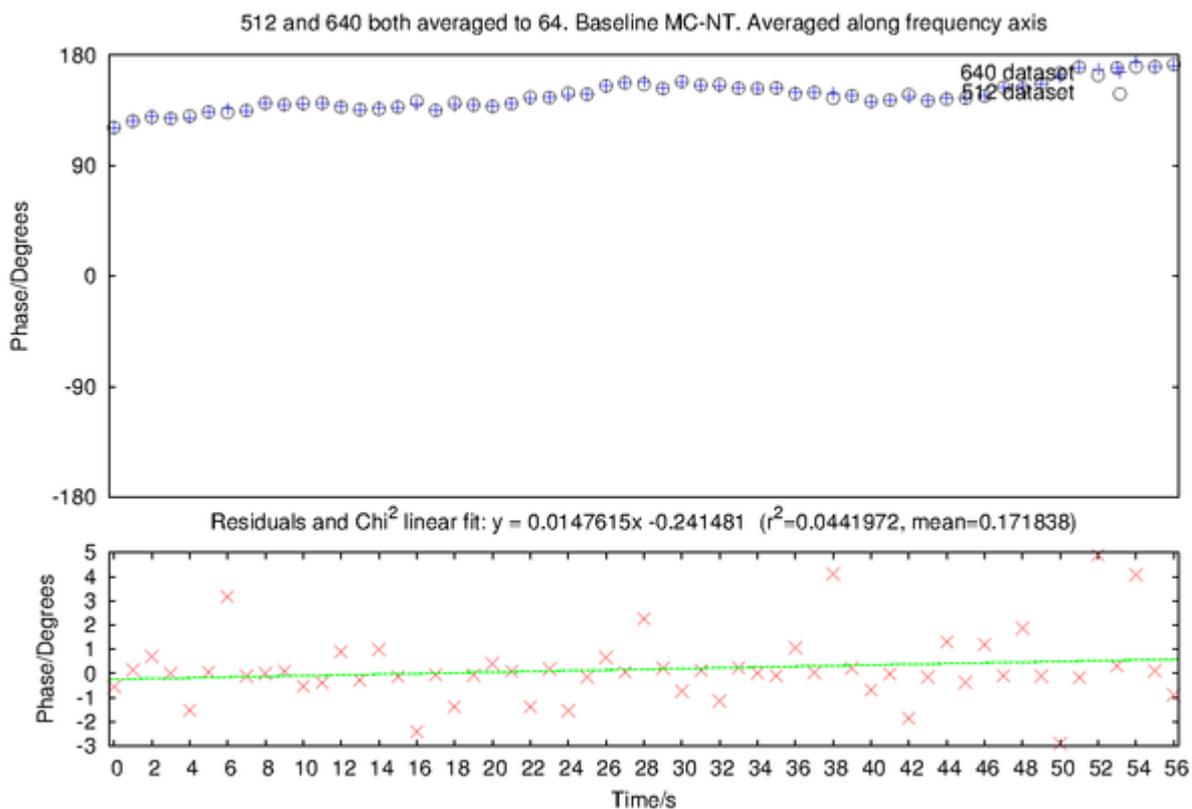
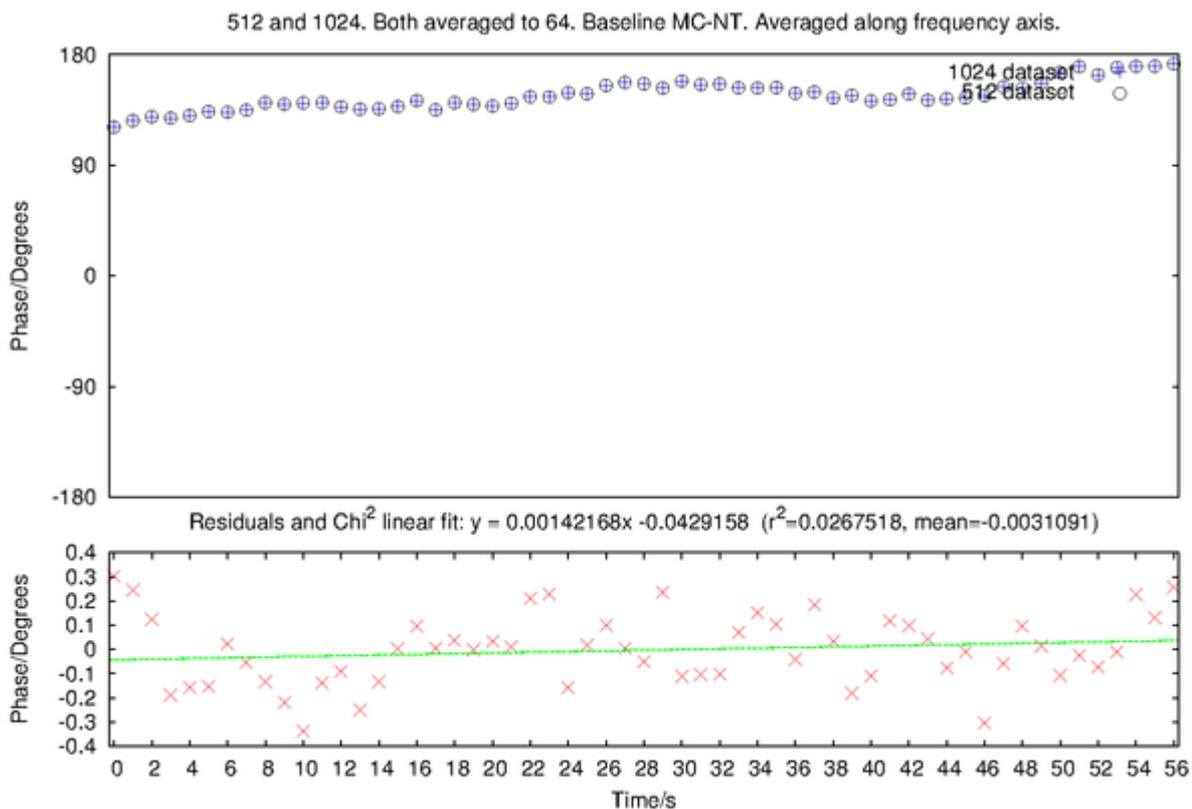
nChan	specAvg	nChanOut	Mean data transfer rate
16	1	16	438.25 Mb/s
500	25	20	281.63 Mb/s
512	32	16	279.21 Mb/s
1000	50	20	239.19 Mb/s
1024	64	16	238.87 Mb/s

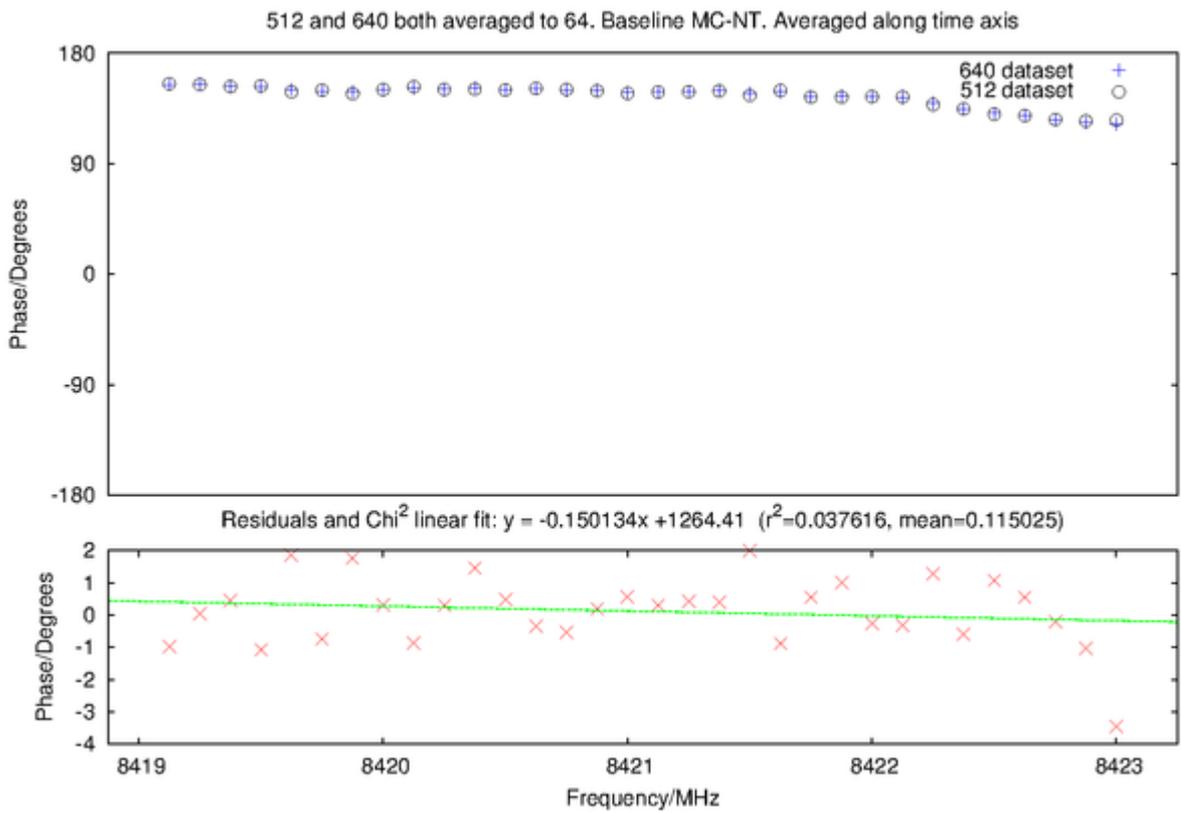
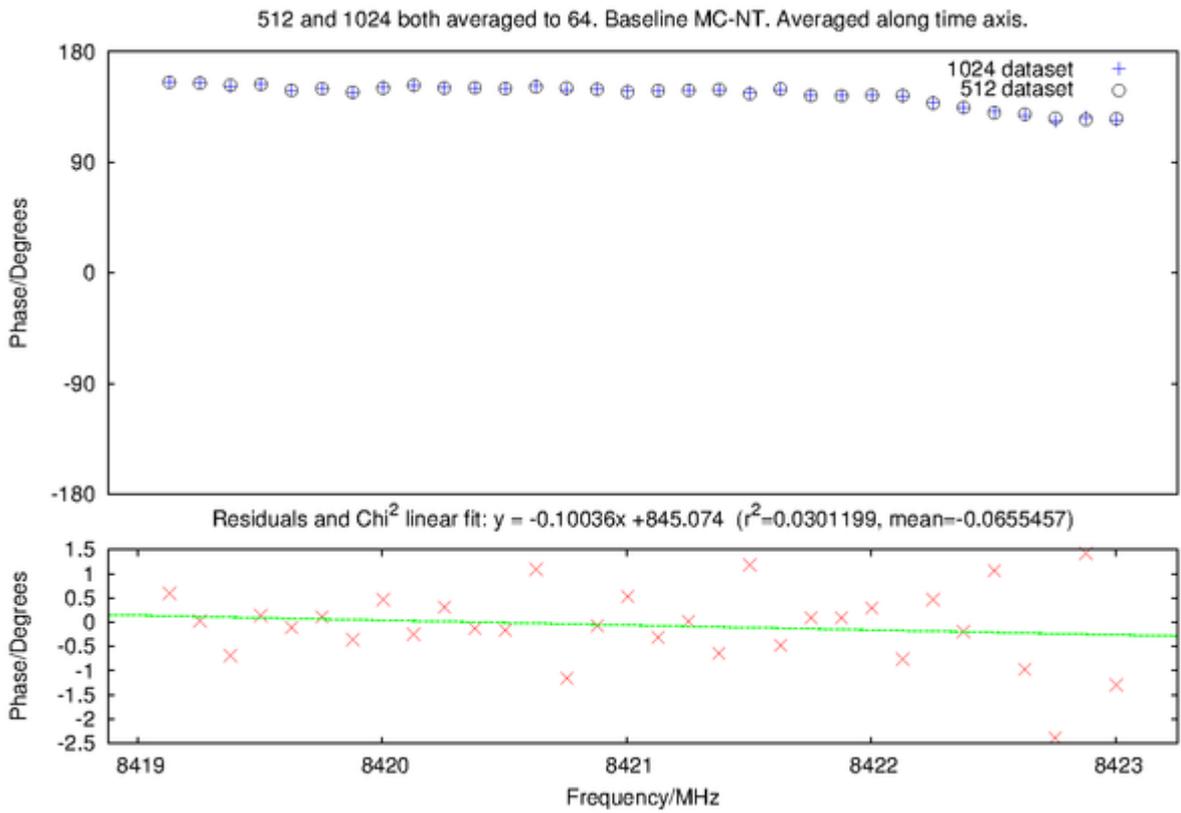
It was not surprising that the prime number was processed the slowest though it was expected that there would be a much sharper drop in performance than was observed. The results also indicate (at least for this many datastreams) that the processing is dominated by the Fourier transforms. It also show that there is no discernible loss in performance between using a DFT or an FFT.

Accuracy

A basic test of accuracy was carried out by comparing the output visibilities from datasets generated different FFT sizes after averaging each down to 64 channels within mpifxcorr (only the central 32 were used for the analysis). After averaging down to the same number of channels, the plots of comparing the phase for 512 vs. 1024 channels are comparable with those for 512 vs. 64 channels. Amplitude is not shown here but there is no significant difference there either.

nFFTChan	nChan	Channels averaged together
512	64	8
1024	64	16
640	64	10





From:

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

Permanent link:

<https://www.atnf.csiro.au/vlbi/dokuwiki/doku.php/difx/difxdft?rev=1309501851>

Last update: **2011/07/01 16:30**

