ALMA Memo 342 - An Improved FX Correlator

John Bunton
CSIRO, Telecommunications and Industrial Physics, Australia
19 December 2000

1. Introduction

The advantages of an FX correlator are well known [Chikada et al, O'Sullivan] but still most correlators use an XF or at best FXF topology. The reasons for this are mainly due to the perceived increase in required data rate and the higher complexity of the FX design. It is also known the advantages of the FX approach increases as the number of inputs to the correlator grows. As ALMA has about 64 antennas [Conway] and the SKA [Smoulders & Van Haarlem] may have up to 1000 antenna stations it is timely to reconsider an FX correlator option. In the following paper, it will be shown that an FX correlator does not necessarily increase either the data rate or bit rate. A design, which achieves this objective, is presented. It is shown that the design, implemented in off the shelf hardware, is competitive in cost with XF correlators built with custom design chips.

2. Caveat

The author apologises if the designs and examples appear to be over-specified but the aim of this paper is to present a design that should have negligible inherent problem. Simplifications or reduced performance may cause errors and the effects of such errors should be modelled for high dynamic-range telescopes such as ALMA and the SKA.

3. Advantages of an FX correlator

In the pure FX[1] correlator the input signal is filtered into a number of frequency bands and a single cross-correlation is calculated for each frequency band. Examples of such designs are the Digital FFT Spectro-Correlator [Chikada et al 1984] and VLBA [Romney] which used an FFT to form the different frequency bands, and the second Westerbork Correlator [O'Sullivan 1984] where analogue filters were used. Some incidental advantages of an FX correlator are the ease with which accurate delay compensation is implemented and the ability to implement a pulsar gate [Romney].

The main advantage of the FX correlator is that dividing the spectrum into a number of frequency bands (say *n* bands) reduces the number of the cross multiply/accumulates[2] (XMACs) required by *n*. To see this consider a hypothetical XF correlator with complex input data at a sample rate *S* and a total useable bandwidth of *S*, (aliasing at the band edge effects ignored). If the final frequency resolution is to be *S/n* then for an XF correlator a total of *k.n* lags of the correlation must be measured. (Where *k*=1 for a rectangular

---

[1] In this paper the frequency domain transformation F is viewed in its most general form and not in the more limited context of the FFT.

[2] Complex data and complex multiplication are assumed throughout.

window and $k \approx 1.7$ for a Blackman window, which may be used to prevent spectral leakage).

Now consider an FX correlator which divides useable bandwidth into *n* equal frequency bands. Each band is *S/n* wide with a sample rate of *S/n*, again ignoring aliasing and band edge effects. A single XMAC is need for each of the n frequency bands. However, the rate at which each of the XMAC units must operate is reduced by a factor of n. The results for the two different correlators is summarised below

	Number frequency band	Data rate per frequency band	Total Data Rate	Lags per frequency band	XMAC ops/sec one baseline
XF correlator	1	S	S	k.n	k.n.S
FX correlator	n	S/n	S	1	S

Table 1 Comparison of XMAC processing for XF and FX correlators

Added to the FX correlator result are operations needed to form the filter bank. An FFT based filterbank requires about $3\log_2(n)$.S operation/sec per antenna. The filterbank outputs then form correlations with themselves and all the other antennas in the array. Thus if the basic FFT operation and an XMAC are comparable then the overhead to implement the filterbank for an array of M antennas is $(3\log_2(n)/M)$.S XMAC operations/sec per correlation. Thus, even for a 6-antenna telescope such as the Australia Telescope (AT) a one million channel FX correlator has a total computational load equivalent to a 11 lag XF correlator, 10 for the filter and 1 for the correlation.

For an XF correlator generating n independent frequency bands the number of XMAC operations can be as high as 1.7*n*. For current telescopes *n* is at least 100 and current specifications for ALMA are for 1024 lags. The theoretical cost advantage of an FX correlator for ALMA is almost 1000 to 1. Of course, current limits on the number of frequency channels have been largely set by hardware limitations. This in turn results from the use of XF correlators. The use of an FX correlator can effectively remove this limitation and wideband correlators with very high frequency resolution become possible.

4. Disadvantages of the FX correlator

4.1. Data rate increases?

The perceived disadvantages of the FX correlator are the cost of the frequency splitting operation, and an increase in the bit and data rate. How the perception of the last disadvantage arose can be seen from a simplistic application of FFT based correlation. After processing two input signals with an FFT, convolution can be obtained with point by point multiplication. But this forms a circular convolution. The problems caused by this are reduced by windowing and overlapping the data [Romney] or a linear convolution can be implemented by setting half the data to zero. Either method increases the output data rate compared to the input data rate, with the rate doubling for linear convolution.

In table 1 it is seen that the data rate to both XF and FX correlators are the same. This is true only if the transition bandwidths as a fractional of total bandwidth are the same for the XF and FX frequency channel. An example of a high quality filter for an FX correlator is given in section 5. Scaled up in frequency this filter would be suitable as the bandlimiting filter in an XF correlator. In both the XF and FX designs the aliasing and band-shape errors would have a similar impact and this occurs with both correlators operating at the same data rate.

4.2. Word length growth?

The second source of increase in data bandwidth is the growth of word length through an FFT. For an input consisting of a sine wave plus noise, say of equal amplitude, the FFT output at the frequency of the sine wave has an amplitude that is much greater than infrequency channels containing just noise. For a 16384-point FFT the amplitude is 128 times larger. This implies that at least 7 extra bits of precision are need to represent the output data. Add to this possible variations in the noise spectrum, and the possibility that the magnitude of the interference could be much larger then the noise, then the output from the FFT may need 14 or more extra bits of precision. In the case where the input data was 2-bit data, the data rate has grown by a factor of 8 or more. It is easy to understand statements seen in the literature that an FX correlator may cause an increase in data rate “by perhaps a factor of 32” [Capallo 2000]. But this results from a simple acceptance that all the output data bits from an FFT are useful. There is also the problem that for 1 and 2-bit input data there has already been loss in sensitivity due to quantising. Any further loss is unacceptable and complicates the Van Vleck correction therefore; the full precision of the FFT should be preserved.

Coarse quantisation should occur after the FFT instead of ahead of it. With properly set quantiser levels the Van Vleck correction is again easily calculated and the input data to the correlator need only be 1 or 2 bit. This is approach taken in LOFAR [Bregman]. The high levels of interference necessitate high-resolution analog-to-digital converter. LOFAR will then use adaptive spectral inference rejection on each of the spectral channels to ensure 2-bit quantisation is sufficient for data going to the correlator.

An alternative to inference rejection is to optimally quantise each frequency channel. After splitting the data into multiple frequency bands, each with a low sidelobe level, the majority of the bands will contain system noise with at most low level interference. For these bands the power in the astronomy signal plus system noise predominates, and when considered in isolation each of these bands can now be digitised by a 1 or 2-bit quantiser. Only those bands that have high levels of interference need multibit quantisation. Assuming 10% of the bands need 8-bit quantisation and another 10% need 4 bits then the average bit rate is 2.8 bits per sample. When a signal with this level of interference is sampled as a wideband signal then six bits of precision are needed to maintain equal signal integrity. The use of optimal quantisation after the frequency splitting operation can actually lead to a significant reduction in bit rate. This phenomenon is used in many modern audio and video data-compression algorithms to reduce the data rate of the encoded signal.

Optimum encoding is achieved by performing a signal analysis operation on each frequency channel. This could be implemented in a similar fashion as current XF correlators. Alternatively, measurement of auto-correlations should provide sufficient data to determine the number of bits and the weighting of each bit for each channel. As this data has a low update rate it is a simple matter to update the digitisers and XMAC units periodically. The multiply operation in the XMAC would use the information to correctly decode the data for each channel. The multiply hardware becomes more complex. Still the XMAC hardware in an FX correlator is orders of magnitude cheaper than that of the equivalent XF correlator.

4.3. Filtering complexity

One definite added cost in an FX correlator is the need to filter the data into multiple bands. But, this cost grows linearly with number of antenna stations as opposed to the XMACs that grow as the square of the number of antenna stations. A DFT filterbank, to be presented in the next section, has a computational load equal to two FFTs, one for the FFT and approximately one for the pre filters. Using this and the results from section 3 it is found that the cost of the filterbanks amortised over all baseline plus the single XMAC per baseline gives:

The equivalent computational load of an FX correlator is equal to that of an XF correlator generating $1+6\log_2(n)/(M-1)$ lags per baseline, where n is the number of frequency channels and M the number of antennas.

This number has to be adjusted for differences in the complexity between FFT operations and XMAC operations in XF and FX correlators. The table gives the equivalent compute load for an FX correlator using the assumption that FX operations is four time the complexity of a XF operation.

	1024 channels	1,000,000 channels
2 antennas	244	484
6 antennas (AT)	52	100
64 antennas (ALMA)	7.75	11.5
300 antennas (SKA?)	4.8	5.6

Table 2 Estimated equivalent computational cost on an FX correlator in units of XF correlator lags

Even for a 2-station 1024-channel correlator the FX correlator is computationally more efficient. For large numbers of antennas, the increasing efficiency means that an FX design implemented in DSP processors and FPGAs is a viable option when compared with a custom chip XF design.

4.4. Other Problems

Chikada 1991 enumerated three other problems with the original FFT FX correlator: variations in the Van Vleck correction due to multiple quantisation, the effects of cyclic convolution and the effects of segmenting the data. Cyclic convolution and segmentation of data problems are eliminated when the output data does not contain aliasing. The Van Vleck correction problem is removed when coarse quantisation occurs only at the output of the filter bank.

4.5. Summary of advantages and disadvantages

It has been shown that an FX correlator need not suffer from a growth in data or bit rate due to the filtering operation. This eliminates data transmission costs as a reason for rejecting the adoption of an FX correlator. If development costs can be contained, particularly by the use off the shelf components, then an FX implementation is the logical choice for future correlators. In the following sections an efficient implementation of a high-quality filterbank and a strawperson hardware design are presented to illustrate the possibilities.

5. Proposal for an FX correlator

The main problem to be solved in implementing an FX correlator is the design of an efficient high-quality filterbank. The FFT has been used traditionally but quality of the filtering implemented has compromised data rate or correlation quality. A standard signal processing method for implementing a high-quality filterbank is shown below, in Figure 1. In this approach each output is derived from the input by first translating the centre of the wanted band of frequencies to zero frequency (multiply by $e^{-j\omega t}$). This is then lowpass filtered and decimated. The filter quality is solely determined by the prototype lowpass filter $h_0(n)$.

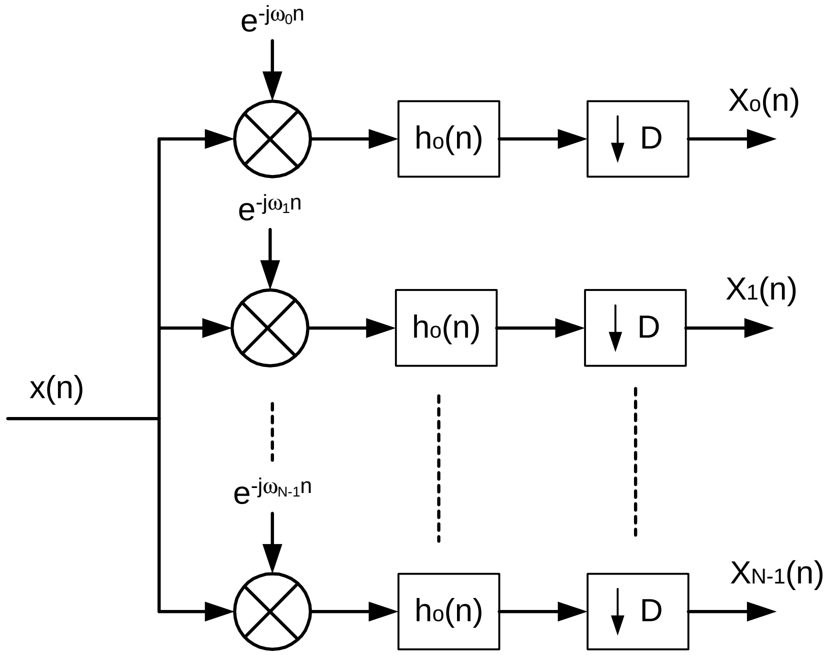

Figure 1 Standard filter bank

If the frequency bands are equi-spaced and the decimation factor D equals the number of bands N then the more efficient DFT filterbank implementation is possible [Bellenger et al, Crochiere& Rader or Proakis et al] as shown in Figure 2. The filters $p_k(m)$ used in the DFT filterbank are a decimated version of the prototype lowpass filter $h_0(n)$.

$$p_k(m) = h_0(Nm + k)$$

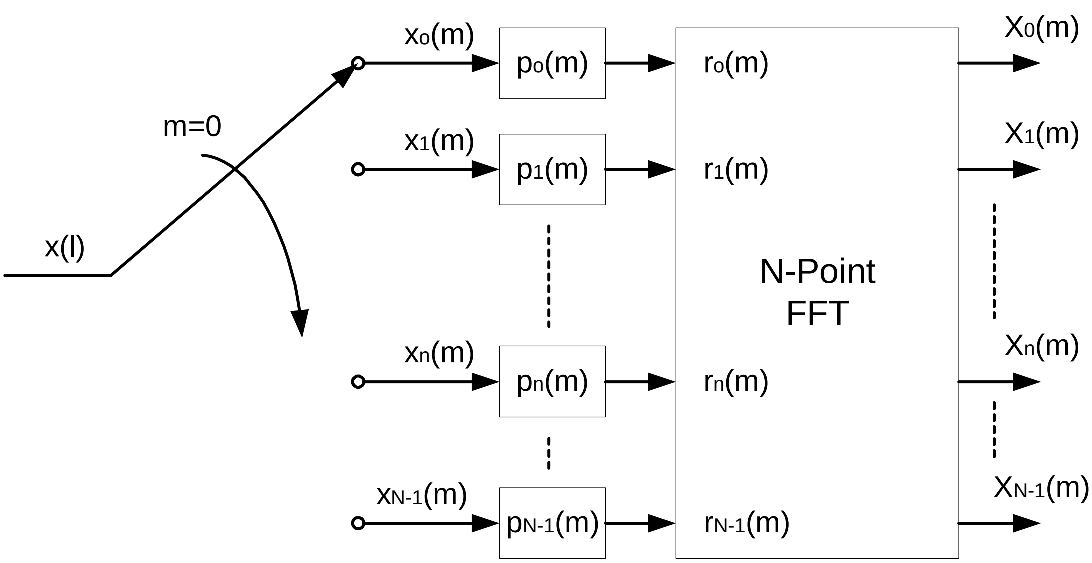

Figure 2 DFT filterbank with arbitrary frequency response

This implementation is the normal FFT approach preceded by a short pre-filter at each input. This has the advantage that the number of operations per sample does not change significantly as the number of channels N increases. The effect of the pre-filter is to convert filtering from a sinc filter, for the straight FFT, to one determined by $h_0(n)$, the prototype lowpass filter. An example of a prototype lowpass filter is shown in Figure 3a. In this example, the filter length is 12 times the length of the FFT making each of the pre-filters a 12-point filter. The filter response achieved when this filter is used to implement the DFT filterbank of Figure 2 is shown in Figure 3b. In this figure, the response of one frequency channel is shown together with the overlapping upper half of one adjacent channel.

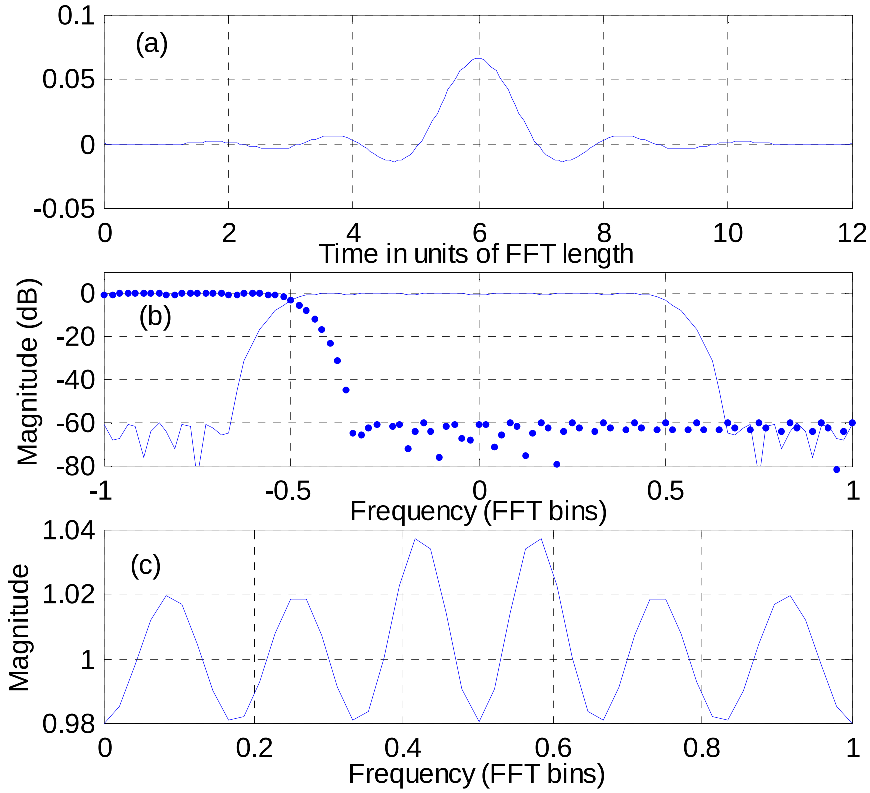

Figure 3 Prototype filter and the response of the DFT filterbank it implements. (a) impulse response of prototype filter $h_0(n)$, (b) magnitude response of one of the outputs, response of adjacent output dotted and (c) sum of the correlation power for the two responses shown in (b)

The accuracy with which a correlator measures the total power of a monochromatic source is shown in figure 3c. The response is the sum of the correlated power in two adjacent frequency channels and assumes aliased frequency components correlate. For the standard FFT approach, the total normalised power for a monochromatic source varies from 1 to 0.82 across the band. For the example of a DFT filterbank shown the ripple is 2% across most of the band. Use of quadrature mirror filters to implement the prototype lowpass filter reduces the error to zero but leads to longer filters. See Romney

1995 for a discussion on the standard solutions for reducing this type of error in current FX correlators.

It is also seen that that the filter in Figure 3 reduces spectral leakage[3] for all except adjacent channels to -60dB or better and after correlation the spectral leakage is below –120dB. The only major imperfection in the system is the aliasing that occurs for a small band of frequencies just above the half channel frequency (the frequency half way between the centre frequencies of two adjacent channels). There is not a problem if no further processing of the signal is needed and the aliases correlate. If the aliases don't correlate then there is a small reduction in sensitivity and some added noise. Aliasing also causes problems if higher spectral resolution or other forms of processing such as interference mitigation are required.

The aliasing is inherent in the implementation shown in Figure 2 because the data rate of any of the FFT outputs is identical to the frequency separation of the channels. Either the frequency response of adjacent channels must overlap, leading to aliasing, or the filter must be made narrower leading to a loss of information at the half channel frequency. If aliasing does occur then the technique that is used in the proposed WIDAR correlator [Carlson & Dewdney] can reduce this problem by decorrelating the aliased signals. With this technique, the correlated aliases tend to zero after integration.

For the filter response shown in Figure 3 the problems of aliasing are eliminated if the output data rate could be increased by a factor of 1.32. This moves the frequency at which aliasing occurs to 0.66 of the FFT channel separation. If aliasing is to be eliminated at frequencies below 0.5 of the FFT channel separation then a sample rate increase of 1.16 is sufficient. A change in sample rate is accomplished by changing the commutation rate and reordering the input data to the pre-filters $p_k(m)$ [Crochiere, Crochiere& Rader]. To see how this can be done consider the diagrammatic representation of the DFT filterbank, Figure 4.

---

[3] Spectral leakage is used to refer to frequency components in the stopband of the filter that are aliased in band when the data is sampled. This term is used to distinguish these components from aliasing due to components within the transition band of the filter.

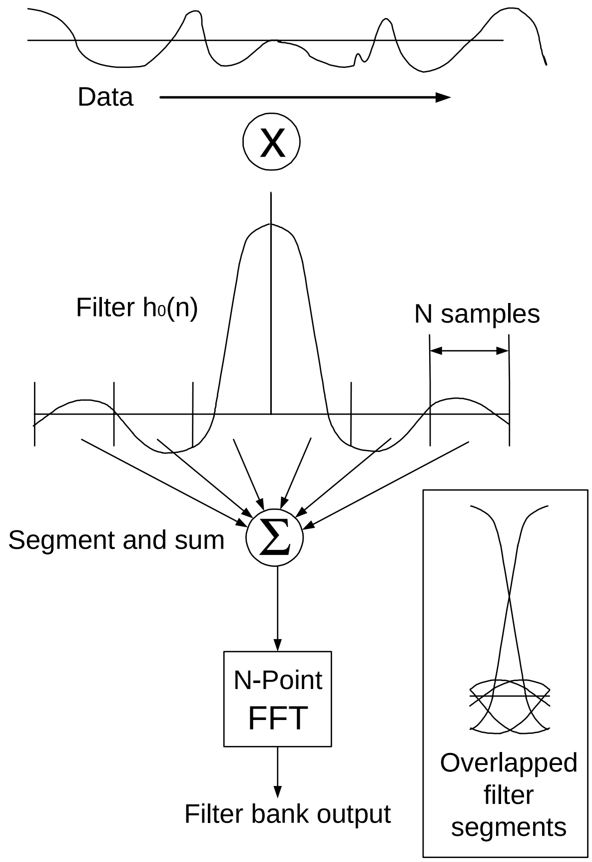

Figure 4 Diagrammatic representation of the operation performed by the system shown in Figure 3.

The input data and the filter $h_0(n)$ is broken up into segments of length N. A point by point multiplication between the data and filter coefficient is performed and the results are summed. The summation is across corresponding elements of each segment. To illustrate this point the segments of the filter are shown overlapped just prior to summation. The summation generates the N samples of input to the FFT. For the implementation shown in Figure 2 the data is shifted by N samples to generate the data for the next input to the FFT. If, instead of shifting N samples, the shift is $N/\boldsymbol{I}$ samples between each FFT then the output data rate is increased by a factor $\boldsymbol{I}$. This can be accomplished by adding a cyclic shift operation to the shift registers implicit in the filter $p_k(m)$. This cyclic shift moves the data for filter i to filter $[i + N/\boldsymbol{I}]_{\mathrm{mod}\ N}$ before each round of input data commutation.

6. A strawperson 2 GHz Correlator

This design is for a 1024 channel FX correlator that uses the prototype filter shown in Figure 3 with an oversampling ratio of 1.16. This sampling ratio allows later stages of processing to achieve higher frequency resolution without aliasing. The design is based upon general purpose DSPs to provide a baseline cost and performance.

The announced, but yet to be released, C64 DSP from Texas Instruments is benchmarked to perform a 1k 16-bit complex FFT in 6002 clock cycles. With expected clock rates of 1.1GHz, FFTs can be executed with a data rate of 183 Megasamples per second. If a 1.16 oversampling ratio is used and an allowance made for data I/O then an actual signal bandwidth of about 125-150 MHz could be supported.

This excludes the pre-filters $p_k(m)$. As the complexity of the pre-filters is of a similar order to that of the 1k FFT a second DSP can be used for this operation. Complex data, rate S, going to the first DSP for pre-filtering is then passed to the second DSP at a rate of $1.16S$ where the FFT is performed.

For new correlator designs bandwidths of 2GHz or more are required. For 2GHz bandwidth this could be implemented by first filtering the data into 125MHz bands and then processing each band with a pair of DSPs. But the DSPs are already performing a band splitting operation. The additional sixteen-channel filterbank is not needed if the input data is processed in blocks. To generate the data for a single FFT, twelve 1024-sample sections of data are needed. To generate the next set of data for the FFT, 883 samples are needed. Thus providing the pre-filter DSP with 883x49 + 1024x12 = 53,105 samples is sufficient to generate the inputs for 50 consecutive FFTs. These FFTs require 51,200 input samples and generate 51,200 outputs, assuming none are discarded. Thus, the input and output data rates are approximately the same for both the pre-filter and FFT DSPs.

A data formatter preceding sixteen DSP pairs is now sufficient to process data at 2 Giga samples per second (2GHz bandwidth with complex data). The formatter sends 53,105-sample blocks of data to each DSP pair in turn with each block 44,150 samples further along the data sequence. The data formatter is significantly cheaper than the sixteen 125 MHz filters it replaces. The only change in system performance is that the channel frequency separation has increased from 122kHz (125MHz/1024) to 1.95MHz (2GHz/1024). At the cost of more memory, the higher frequency resolution can be recovered by increasing the length of the FFT. This results in a small decrease in system bandwidth due to the increased FFT processing.

7. Higher spectral resolution, data reformatting and optimal bit assignment

In general, high frequency resolution is needed in only a small part of the spectrum. This can be added at a small incremental cost with 1 or 2 extra DSPs. The important element in achieving this is the reformatting of the data after FFT operation. Each FFT DSP generates output data for 1/16 of the total time. The sixteen time-division-multiplexed data streams must be reformatted into 1024 continuous frequency channels. Each frequency channel has a data rate of 2.266 MHz (1.95MHz times the oversampling ratio), so in a practical system the channels might be aggregated into groups of 64 to 256 channels. In an intelligent reformatter there would be complete freedom to choose which frequency channels are assigned to each group. One or two groups could then be further processed by a DFT filterbank to provide high-resolution spectroscopy on a group of arbitrarily chosen 2 MHz bands. A single extra DSP should be sufficient to process up to 32 individual frequency channels.

The ability to reformat the data also provides a simple means for reducing the bit rate. Because of bus bandwidth limitations, each group of channels will have the same maximum bit rate. This bit rate can be assigned in many different ways. For example, a 1024 Mbit/s bus can support 256 2-bit complex channels or 64 8-bit complex channels. To ease the implementation of this the FFT DSP could save data in 32-bit words where one word could have 2 8-bit complex, 4 4-bit complex or 8 2-bit complex values stored. Six of these busses might be used to provide, for example, three 256 by 2-bit, a 128 by 4-bit and two 64 by 8-bit channels for a total of 1024 channels. With these groupings, there is an average of 3 bits per sample (6 per complex value). The higher precision channels are only needed where there is significant interference or the data is not noise like. The only drawback of this adjustable bit rate scheme is that the XMACs need to handle the different precisions. This is not as costly as it might first appear because an 8-bit multiply has a complexity equivalent to two 32-bit accumulators making the FX XMAC equal in complexity to three XF XMACs. The comparison shown in table 2 assumes the complexity is four times greater.

A final mode of operation that could be considered for the system is to sub divide the system into a number of lower bandwidth sections. This is achieved by including extra modes in the data formatter and reformatter so that a number of lower sample rate signals are processed by subsets of the DSP pairs. The reprogrammability of the DSP allows many different spectral resolutions to be implemented. For instance, a single DSP pair operating on 100MHz data could provide 100Hz-resolution aliasing channels with a 1Meg FFT. Reprogrammability can also be used at the full 2GHz rate to vary the number of frequency channels as long as there is sufficient buffer memory in the formatter and reformatter,

8. Data and Bit Rate for the Strawperson Design

If the entire spectrum is preserved, without aliasing, at the output of the DFT filterbank then there is an increase in data rate by a factor of 1.16. But in any practical design the data at the band edges may not be useful due to low amplitude in the transition band and aliasing. Up to 10% of the spectrum may be degraded in this manner. If this section of the spectrum is discarded there is little loss of sensitivity. In fact, discarding this part of the spectrum and reusing the bus bandwidth to provide extra bits of precision in other parts of the spectrum can lead to a higher sensitivity system. With 10% of the spectrum discarded, the data rate increase at the output of the DFT filterbank is 5% and this is for a system where there is no aliasing across the span of each frequency channel.

Samplers on radiotelescopes have traditionally been 2, 3 or 4 level samplers. With the increasing presence of interference as bandwidth and human activity increase it will be necessary to increase the number of sampling levels in future designs. The samplers will generate 3 or more bits per real sample. Consider a single 2MHz band where the power level of an interferer is 10 times the total system noise. Such an interferer has a signal level 3 times higher than the system noise. Just under two extra bits of sampler precision are needed to maintain the difference between sampler levels at the same value as the no interferer case. If there were 10 2MHz channels with this level of interference then over the full 2GHz the average increase in data rate is 0.02 bits. When measured relative to the full 2GHz band the total power in the 10 interferers is 3 times higher than the system noise power. (Total interferer power is 100 times and total system noise is 32 times greater than the 2MHz system noise.). In an XF correlator a full extra bit of sampler precision is needed to maintain signal qualtity. In this case, the use of a DFT filterbank with optimal bit assignment to the data leads to a reduction in the bit rate when compared to the bit rate needed for an XF correlator.

As stated previously, the requantisation at the output of the DFT filterbank means it is prudent not to degrade the input data by the use of overly coarse A/D quantisation. Any data loss at the A/D quantiser cannot be recovered and the FFT quantiser can degrade the data further. Thus with a DFT filterbank it is best to sample and process the data at a high precision and apply any coarse quantisation at the output of the DFT filterbank.

9. Cost comparison

A current XF correlator chip runs at 128 MHz and can process 1024 real lags. A newer design should be able to process 4096 real lags at 125 MHz [Escoffier]; call the cost of this chip $X. The FX correlator XMACs could be implemented using XILINX XCV812E devices which contain 1120kbits of RAM. This RAM is sufficient to implement 23k complex (equivalent to 47k real) FX correlations at 24-bit accumulator precision. It is estimated that when in production a custom XF chip will be comparable in price to the XILINX device. This XILINX device implements over 10 times as many frequency channels of correlation as the custom chip. Alternatively, using a high pin-count FPGA and external memory (12 devices 64kx16-bit) produces a module that forms cross correlation for four baselines with 65536 complex frequency channels per baseline. This module would cost less than the estimated cost of the custom XF chip and produce 128 times as many frequency channels.

For a 125 MHz span of the 1024-channel DFT filterbank the cost for each antenna is estimated to be equal to 2.5 $X. One XF custom chip cost for each of the pre filter and FFT DSPs and a quarter of an XF chip cost for the data formatter and reformatter. The data formatter and reformatter are comparatively simple and will probably be implemented in an FPGA.

For M antennas the cost of custom chips for a N by 4096lag 125 MHz XF correlator is approximately $X.N.M.(M-1)/2. For the FX correlator the same frequency resolution is obtained with a 2048N-channel filterbank. The approximate cost of this filterbank $\$X.2.5(\log_2(2048N)/\log_2(1024))$ per antenna. Assuming a cost of $X.M.(M-1)/20 for the XMACs in the FX correlator then the system costs are equal when

$$2.5.([11 + \log_2(N)]/10)M + (M-1)M/20 \qquad = N(M-1)M/2$$

$$M \quad = 5.\,[11 + \log_2(N)]/[10N - 1] + 1$$

Thus, estimate costs are equal for a 7-antenna array with a 4096 lag or 2048 frequency channel correlator. These costs are very rough estimates but they do show that even for a small number of antennas an FX correlator using off-the-shelf components can be competitive with an XF correlator using a custom design chip.

10. Conclusion

An FX correlator design has been demonstrated that has none of the defects inherent is some of the previous design. Even for radiotelescopes with a small number of antennas, this FX correlator implemented in DSPs and FPPGAs is cost competitive with a custom chip XF correlator. For a radiotelescope such as ALMA an order of magnitude cost reduction is possible.

For this to be achieved the following changes to current FX correlators designs are needed:
1. Improved filtering possibly based on the DFT filterbank
2. Higher A/D precision to allow
3. Coarse quantisation at the output of filterbank with level control similar to that currently implemented with XF correlators.

The first improvements ensure there is no or negligible increase in the number of data samples going to the correlator and the final two ensure there is no increase in the number of bits per sample

11. Acknowledgement

The author wishes to thank Colin Jacka and Dick Ferris for their careful critique of this paper and latter also for many interesting discussions on correlators and signal processing and in particular for bringing the DFT filterbank to the authors attention.

12. References

Bellenger, M.G., Bonnerot, G., & Coudreuse, 'Digital Filtering by Polyphase Network: Application to Sample Rate Alternation and Filter Banks', IEEE Trans. Acoustics Speech and Signal Processing, Vol ASSP-24, No 2, pp. 109-114, April 1976

Bregman, J.D., 'Design concepts for a Sky Noise Limited Low Frequency Array', Smolders, A.B. and Van Haarlem M.P (eds), Perspectives on Radio Astronomy: Technologies for Large Antenna Arrays, 'Netherlands Foundation for Research in Astronomy - 1999

Capallo, R., 'Large-N Correlators for the SKA' Square Kilometer Array US Consortium, Arecibo Observatory, Feb 28-29, 2000. Available at http://www.naic.edu/~ska/usskameeting.htm under Correlators [Roger Cappallo…]

Carlson, B.R., Dewdney, P.E., 'Efficient wideband digital correlation', Electronics Letters, IEE, Vol. 36 No. 11, p987, 25th May, 2000

Chikada Y., et al., ‘Digital FFT Spectro-Correlator for Radio astronomy’, in Indirect Imaging (ed J.A. Roberts), Cambridge University Press, pp387-404, 1984

Chikada Y., ‘Correlators for Interferometry – Today and Tomorrow’, present in Astronomical Society of the Pacific Conference Series, Vol 19, ‘Radio Interferometry: Theory, Techniques, and Applications’ (eds Cornwell, T.J. and Perley, R.A.), 1991

Conway, J., ‘Observing Efficiency of a Strawperson Zoom Array’, ALMA Memo #283, 2000, available at http://www.alma.nrao.edu/memos/index.html

Crochiere, R.E., ‘A Weighted Overlap-Add Metrhod of Short-Time Fourier Analysis/Synthesis’, IEEE Trans. Acoust. Speech Signal Process., Vol ASSP-28, No 1, pp 99-102, Feb 1980.

Crochiere, R.E. & Rader, L.R., ‘Multirate Digital Signal Processing’, Prentice-Hall, New Jersey, 1983

O’Sullivan, J.D., ‘The Westerbork Broadband Continuum Correlator System’, in Indirect Imaging (ed J.A. Roberts), Cambridge University Press 1984, pp405-414

Proakis, J.G., Rader, C.M., Ling, F. & Nikias, C.L., ‘Advanced Digital Signal Processing’ Maxwell Macmillan 1992

Romney, J.D., ‘Theory of correlation in VLBI’ present in Astronomical Society of the Pacific Conference Series , Vol 82, ‘Very Long Baseline Interferometry and the VLBA’ (eds Zensus, J.A., Diamond, P.J. & Napier P.J.) 1995

Smolders, A.B. and Van Haarlem M.P (eds), Perspectives on Radio Astronomy: Technologies for Large Antenna Arrays, ‘Netherlands Foundation for Research in Astronomy - 1999