Introduction to the Spectral-Domain ("FX") Correlator

JONATHAN D. ROMNEY
National Radio Astronomy Observatory
Charlottesville, Virginia

1986 May 5

To begin at the end: what we want at the output of the VLBA correlator is, in all cases, the cross-power spectrum. This assertion is obvious in the case of spectroscopic observations and, with the introduction of bandwidth synthesis and global fringe-search techniques for continuum measurements, has come to apply universally. It is therefore attractive to consider exploiting the Fast Fourier Transform (FFT) algorithm to perform the necessary frequency analysis on individual station data streams before — rather than on a baseline basis after — "correlation". The latter term is something of a misnomer in this case, since the cross-correlation function need never actually be evaluated, but nevertheless we will term this route to the cross-power spectrum "correlation in the spectral domain". The well-known convolution theorem of the Fourier transform ensures that pairwise multiplication of the Fourier-analyzed station spectra is equivalent to applying the transform to the conventionally formed baseline lag correlation functions.

An elegant depiction of these two alternative data paths, with an additional dimension imposed by the aperture – image transform relationship, is given by Chikada *et al.* (URSI/IAU Symposium on Measurement and Processing for Indirect Imaging, ed. J. A. Roberts, 1983, p. 387), who describe the implementation of a spectral-domain correlator for the millimeter interferometer at Nobeyama Observatory. Those authors also introduce the convenient "FX" nomenclature which we will adopt: F refers to the Fourier transform operation, while X is multiplication or (in the conventional "XF" lag scheme) cross-correlation; the order of operations is simply read left-to-right (not the opposite as in mathematical operator notation!).

This memorandum contrasts the two correlator concepts, outlining the salient advantages which, we believe, recommend the FX architecture as appropriate for the VLBA correlator. Economies in hardware requirements are emphasized, and illustrated by a simple heuristic calculation; a more detailed analysis will have to be deferred until completion of detailed cost estimates for the FX correlator, which we are confident will compare very favorably with those for the XF system proposed in Correlator Memo 41 and subsequent addenda. Beyond these considerations, correlation in the spectral domain entails a variety of technical benefits which are also described briefly. A final section summarizes the known disadvantages of, and a few of the many unanswered questions about the FX approach.

Hardware Economies

The principal advantage of the FX architecture is a significant reduction in the multiplier/accumulator hardware required. This is easily demonstrated by considering the aggregate multiplier rates required for the spectral and lag correlators. Both schemes can be characterized by the same set of essential parameters:

	n_s	Number of stations;
	n_c	Number of channels ("IFs") per station;
	n_t	Number of time or lag points per Fourier transform;
and	r_0	Input sample rate.

Assume that the spectral correlator's FFTs are implemented using a straightforward radix-2 Cooley-Tukey algorithm; more efficient methods using, *e.g.*, radix-4 or -8 stages are attractive possibilities but this will not alter the results of the present calculation. Then the transform consists of $\frac{1}{2}n_t log_2 n_t$ "butterfly" operations of 4 multiplies each, executed at a rate r_0/n_t. The aggregate multiplier rate for station/channel FFT operations in the FX correlator is

$$r_{F(X)} = 2r_0 n_s n_c log_2 n_t. \tag{1}$$

These station spectra are combined by pairwise multiplication to form $n_s^2/2$ baseline cross-power spectra (including the real single-dish "auto" spectra), again at the transform rate r_0/n_t. Only $n_t/2$ points in each station/channel spectrum contain signal power, and since the spectra are complex, 4 multiplies are required. The aggregate cross-multiplication rate in the FX case is then

$$r_{(F)X} = r_0 n_s^2 n_c, \tag{2}$$

and for the combined FX process

$$r_{FX} \equiv r_{F(X)} + r_{(F)X} = r_0 n_s n_c (2log_2 n_t + n_s). \tag{3}$$

For comparison, the lag correlator forms the same number of baselines, each of n_t lags, at the full input sample rate r_0. Assume that fringe rotation is performed in the correlator, so that each lag has one real and one complex input, and thus requires 2 multiplies. Thus the aggregate multiplier rate for the cross-correlation is

$$r_{X(F)} = r_0 n_s^2 n_c n_t. \tag{4}$$

The contribution of the subsequent FFT operation can be neglected, since these lags can be accumulated for some time beforehand:

$$r_{(X)F} = 0; \tag{5}$$

for the combined XF process

$$r_{XF} \equiv r_{X(F)} + r_{(X)F} = r_0 n_s^2 n_c n_t. \tag{6}$$

While only multiplies have been considered in this analysis, to a very good approximation an addition accompanies each multiplication in all cases. One oversimplification must be acknowledged, however: the multiplies are of different complexities, with the XF scheme having the benefit of 1-bit multiplies while most of the FX multiplies are considerably more complex. Chikada *et al.* attempt to take account of this by including a "complexity factor" (which is, unfortunately, quite technology-dependent) in their analysis.

With this caveat, inspection of the equations above shows the superiority of the FX architecture when n_t is large. The extra factor of n_t in eq. (6) compared to (3) is the gain achieved by forming the spectra on a station basis. The F stage conserves the sample rate, producing $n_t/2$ complex output values for each n_t real input samples. (The outputs require more bits, however, so this stage is actually a data expansion operation!) But each of these values must be combined with only *one* other, per baseline, in the X stage. Compare this to cross-correlation in the XF scheme, where each sample is combined with n_t others per baseline.

Eq. (3) shows that the investment required to achieve this gain is modest. For the VLBA, where we would like to support $n_t = 2^{10}$ and $n_s = 20$, both quantities summed in parentheses in eq. (3) are equal, so that

$$r_{FX}/r_{XF} = 2/n_t. \tag{7}$$

Comparison of eqs. (3) and (6) also illustrates how the arithmetic imperatives governing trade-offs among n_s, n_c, and n_t differ between the two schemes. The influence of n_t, in particular, is much diluted by the logarithmic dependence in the FX architecture, inducing at most a factor of two difference in r_{FX}. Thus the three "major modes" of the XF system described in Correlator Memo 41, which trade off n_t against n_s^2, cannot be distinguished in the FX correlator. This is as much an expedient as it is a liability, because the switching required to support these modes contributed much to the complexity of that design.

OTHER ADVANTAGES

In addition to the hardware economies achievable under the FX approach, a variety of technical advantages ensue from organizing the processing in an ultimately station-based manner, with the data available in both time and frequency domains prior to correlation.

Fringe Rotation. Implementing station-based fringe rotation in a lag correlator requires the transmission and correlation of multi-bit rotated samples to avoid loss of sensitivity and spurious correlation at harmonics of the correct station phase rates, and this has proven impractical. The FX scheme, in contrast, must already produce multi-bit samples at the output of the F stage, and it is a minor additional burden to start the transform with fringe-rotated, complex values of sufficient precision to preclude the two problems mentioned.

Closure Errors. The FX correlator should be less vulnerable to baseline-dependent systematic effects, since almost all of the signal processing is performed on a station basis. An XF system imposes a diffused, ubiquitous software burden of ensuring that operations on separate baselines are done consistently.

Fractional-Sample Error. Delay-tracking error can be as large as half a sample period with Nyquist-sampled data, which produces phase errors up to a quarter-turn at the band edge. Efficient correction for this effect must be applied in the frequency domain, and can easily be implemented as an extra step following the last FFT "butterfly" in the FX correlator. The XF approach, however, is faced with the much larger task of applying the correction on a baseline basis after transforming; in a practical large system only a statistical correction is feasible, which yields the proper amplitude spectrum but sacrifices sensitivity.

Modularity and Expandability. The functions of the two distinct fundamental stages are clearly separable in the FX architecture, and the baseline-dependent X stage is a small fraction of the corresponding XF system and can be designed initially to accommodate reasonable expansion plans. These features facilitate construction and testing in sections if necessary, and allow expansion in both stations and channels to be foreseen.

Adaptability. Extension of the specifications to include extraordinary "stations" — in particular an orbiting interferometer element — is just a matter of adding a special station processor to an existing FX system. Any peculiarities of delay and phase tracking are localized in this processor, and do not pervade the entire correlator.

Disadvantages and Unresolved Questions

A few areas have been identified in which the FX architecture compares unfavorably with the conventional lag correlator, although it may be possible to devise methods for dealing with some of these problems satisfactorily. Therefore they are listed here together with a number of other questions for which we currently have no clear answer.

Invalid Data. There is (evidently) no feasible method of interpolating in real time a best value for data samples which are corrupted by instrumental effects. A purist would maintain that an entire FFT segment of perhaps 1024 samples is thereby invalidated by one missing datum, although realistically one sample forced to an arbitrary value clearly will have only a minor effect. An FX correlator for the VLBA will probably just maintain a count of valid samples and apply a user-designated selection criterion before accumulation. The XF scheme is superior in this regard, since known invalid samples can simply be omitted from the correlation.

Sensitivity. The segmentation of input data inherent in the FX concept also imposes a penalty in sensitivity. Since the "boxcar" windows are applied prior to correlation, the spectral weighting function has a $[\sin(\pi f)/\pi f]^2$ form instead of the $\sin(\pi f)/\pi f$ characteristic of the XF correlator. Several corrective measures are possible, principally interleaved overlapping segmentation and overresolution with subsequent averaging in frequency, but the costs and side-effects have not yet been analyzed completely.

Quantization. Corrections for quantization in the sampling process are (probably) most easily accomplished in the lag domain. We do not have a proper correction algorithm for the FX scheme at present, and this area is the major open question at present.

Precision. The butterfly processors of the FFT hardware must accommodate a wide dynamic range of spectral forms, from continuum to extremely strong and sharp emission lines. The requirements for these two extremes must still be elaborated. It may be possible to design butterflies which operate with different scaling according to an externally specified mode.

Fringe Rotation. The expedience of incorporating fringe rotation into the F stage was cited above as a prime advantage of the FX correlator. However, optimization of the form of the rotator function has not yet been completed. This issue includes both the precision with which the phase is supplied and the choice of levels in the function.