Astronomy & Astrophysics

A spectral synthesis method to suppress aliasing and calibrate for delay errors in Fourier transform correlators

T. Kaneko and K. Grainge

Cavendish Laboratory, Cambridge University, Cambridge CB3 0HE, UK
e-mail: tk229@mrao.cam.ac.uk

ABSTRACT

Context. Fourier transform (or lag) correlators in radio interferometers can serve as an efficient means of synthesising spectral channels. However aliasing corrupts the edge channels so they usually have to be excluded from the data set. In systems with around 10 channels, the loss in sensitivity can be significant. In addition, the low level of residual aliasing in the remaining channels may cause systematic errors. Moreover, delay errors have been widely reported in implementations of broadband analogue correlators and simulations have shown that delay errors exasperate the effects of aliasing.

Aims. We describe a software-based approach that suppresses aliasing by oversampling the cross-correlation function. This method can be applied to interferometers with individually-tracking antennas equipped with a discrete path compensator system. It is based on the well-known property of interferometers where the drift scan response is the Fourier transform of the source's band-limited spectrum.

Methods. In this paper, we simulate a single baseline interferometer, both for a real and a complex correlator. Fringe-rotation usually compensates for the phase of the fringes to bring the phase centre in line with the tracking centre. Instead, a modified fringe-rotation is applied. This enables an oversampled cross-correlation function to be reconstructed by gathering successive time samples.

Results. Simulations show that the oversampling method can synthesise the cross-power spectrum while avoiding aliasing and works robustly in the presence of noise. An important side benefit is that it naturally accounts for delay errors in the correlator and the resulting spectral channels are regularly gridded

Key words. instrumentation: interferometers – techniques: interferometric – techniques: spectroscopic

1. Introduction

The observing band of radio interferometers frequently needs to be split into sub-bands, either for spectral line observations or to reduce the effects of chromatic aberration. Fourier transform correlators offer an efficient method for dividing the observation band. In this scheme, the signal from the two arms of the interferometer are correlated at discrete delay steps, making direct measurements of the cross-correlation function (see Fig. 1). Taking the Fourier transform of the cross-correlation function gives the complex cross-power spectrum. For a signal of bandwidth $\Delta\nu$, Nyquist sampling theorem requires the signal to be sampled at time intervals of $1/(2\Delta\nu)$ to avoid aliasing. But the cross-correlation function of a band-limited signal extends over an infinite delay range. So the Nyquist sampling theorem holds true only if we sample the cross-correlation function over an infinite delay range. Clearly this is not practical so the cross-correlation function is sampled over a finite range. This results in a recovered signal spectrum with tapered band edges which will overlap with its images in the spectral domain. This overlap causes aliasing and will corrupt the recovered cross-power spectrum.

As an illustration, the plots in Fig. 2 show snapshots of the cross-correlation functions and the recovered spectra. The left-most plots (case 1) show the cross-correlation function when the source transits an east-west baseline (that is when the path difference between the two arms of the interferometer is zero). The circles represent the measurements at discrete delay steps of the correlator. The next plots to the right (case 2) show the cross-correlation function a short while later. The right-most plots show the spectrum calculated by taking the Fourier transform of the discrete measurements. The filled grey circles indicate the spectrum at transit (case 1) and the dark circles are for case 2. In the top row of Fig. 2a, the cross-correlation function was sampled at 16 delays. This critically samples the $\Delta\nu = 6$ GHz signal bandwidth but undersamples the 0−12 GHz basebandsignal, so the positive and negative halves of the 6−12 GHz bands lie side-by-side. In case 1, the cross-correlation measurements trace out a delta function and give the characteristic flat spectrum. But in case 2, the amplitudes are perturbed. Ideally, the amplitudes should not change so this points to a fundamental problem. This could be worked round by oversampling the baseband signal at 64 or more delays as illustrated by Fig. 2b. Now the two sets of amplitudes at different times match perfectly. Oversampling introduces a buffer between the two halves of the passband and also between their spectral images. This suppresses aliasing by reducing the overlaps between the signal and image bands. But clearly, sampling at 64 delays in hardware is not practical.

Aliasing is manifested in the Fourier Transformed data (the spectrum) as temporal modulations in both the amplitude and phase (see Fig. 3). In aliased signals, the noise components between the channels will also be correlated so the individual channels cannot be strictly treated as independent measurements. These effects are particularly pronounced in the edge frequency channels as seen in Fig. 3 so these channels are usually rejected. This may be acceptable in systems with tens to hundreds of
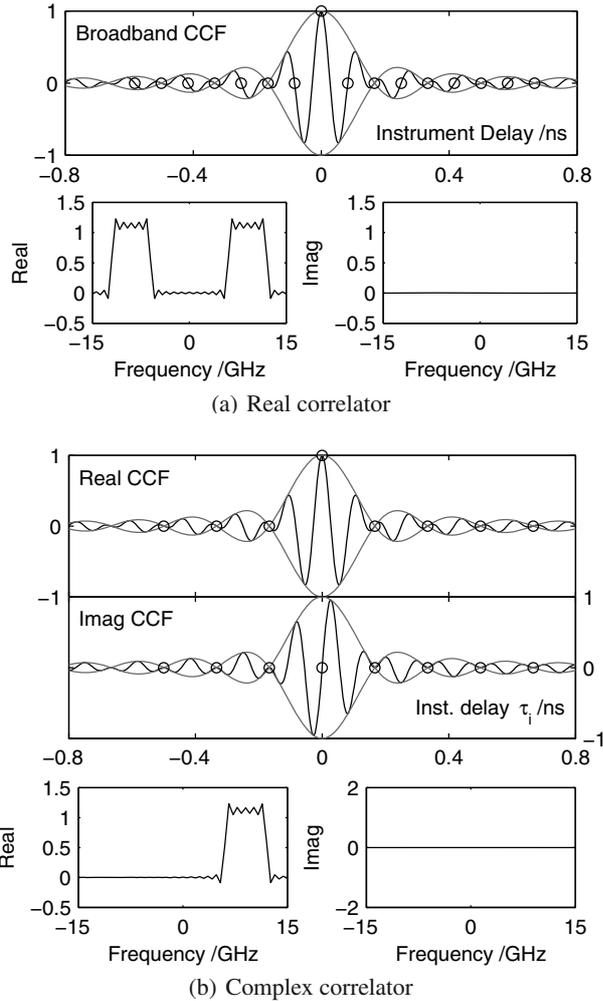
(a) Real correlator

(b) Complex correlator

Fig. 1. a) *Top plot*: the cross-correlation function measured by a real correlator (black) and the envelope (grey). This is a graphical representation of the cross-correlation function in Eq. (1) with the geometrical delay $\tau_g = 0$. The open circles indicate the position of the detectors in instrument delay τ_i. *Bottom plots*: the cross-power spectrum recovered by applying the Fourier transform on the cross-correlation function. We have assumed the telescope parameters in Table 1. **b)** The in-phase (real) and quadrature (imaginary) cross-correlation functions sampled by a complex correlator. The cross-power spectrum recovered in this case is single-sided.

channels. But some correlators, particularly broadband analogue Fourier transform correlators with channels of order 10 (for example, Li et al. 2004; Roberts et al. 2007; Holler et al. 2007, hereafter H07), the loss constitutes a significant portion of the total bandwidth.

In principle, there are three ways to suppress aliasing:

1. increasing the delay range over which the cross-correlation function is sampled. The spectral channel width will be narrower and the edges of the passband will be sharper. So the overlap between the signal spectrum and its images will be narrower. The edge channels will still have to be rejected but it will be a smaller portion of the whole data;
2. oversampling the cross-correlation function at finer delay steps as already illustrated with Fig. 2;
3. reducing the bandwidth so that it is not critically-sampled.

Implementing either of the first two modifications in hardware would be costly and the third approach would lose sensitivity. In Sect. 3, we propose a software-based approach to oversample the cross-correlation function and avoid aliasing. This is based on the well-known notion that the spectrum of a source can be obtained from the cross-correlation function measured by a drift scan. We can reconstruct the cross-correlation function using a combination of source tracking, path compensation and fringe-rotation. In Sect. 4, we will illustrate an application of this technique with simulations. The system requirements are: (1) individually-tracking antennas and (2) discrete-delay path compensation. We have modelled an analogue correlator here but the principles could equally apply to digital correlators.

A number of groups recently reported broadband analogue Fourier transform correlators suffering from delay errors (for example, Harris & Zmuidzinas 2001; H07; Roberts et al. 2007). Simulations by H07 showed that delay errors make the effects of aliasing worse. Although the delay errors can be calibrated a sample at a time (recently by Harris & Zmuidzinas 2001; H07), the method described here will be a natural way of accounting for delay errors. This side benefit is perhaps as significant as the alias suppression aspect of the method. The resulting spectral channels are also regularly gridded at the desired frequencies.

However there are a number of issues that must be considered in a practical system and these are discussed in Sect. 4.3. This paper is a follow up to H07 where we described the development of a broadband 6−12 GHz Fourier transform correlator for the Arcminute Microkelvin Imager (AMI, Kaneko 2006; Zwart et al. 2008). AMI is a new interferometer designed to survey for clusters of galaxies by exploiting the Sunyaev-Zel'dovich effect (Sunyaev & Zel'dovich 1972). The effects of aliasing in the correlator used in this instrument are discussed further in H07. Before going into a detailed discussion of the oversampling method, we will first give a brief overview of how the data is processed in a conventional interferometric system.

2. Conventional interferometry

As the antennas of an interferometer track a source, the geometric path difference (τ_g) between the two arms needs to be corrected (Fig. 4). One approach called path compensation is to insert discrete lengths of delay lines (τ_{pc}). Usually, the signal at the observation frequency ν_{RF} is downconverted by mixing it with the local oscillator (LO) at ν_{LO}. This converts the signal to a lower intermediate frequency (IF) at ν_{IF}. The path compensation is usually inserted in the IF and we will assume that this is the case. It is possible to eliminate the path compensation systems all together by co-mounting the antennas. But our spectral synthesis method cannot be applied to such a system and must have individually-tracking antennas. Individually-tracking atennas also give better control on systematic errors because non-astronomical noise and other contaminating signals can be removed by fringe-filtering.

An interferometer is a device that measures the cross-correlation function of the signals received by a pair of antennas. For a source at the tracking centre, this is given by

$$R_1 = R_0 \frac{\sin \pi \Delta\nu(\tau_g - \tau_i)}{\pi \Delta\nu(\tau_g - \tau_i)} \cos\left[2\pi\left(\nu_{LO}\tau_g \pm \nu_{IF}(\tau_g - \tau_i) + \phi_{LO}\right)\right]. \quad (1)$$

The coefficient R_0 is proportional to the flux density of the point source and also includes factors like the gain of the detector circuit. The plus sign in the ±term applies to upper side-band systems and the minus sign is for lower side-band systems. ϕ_{LO} is
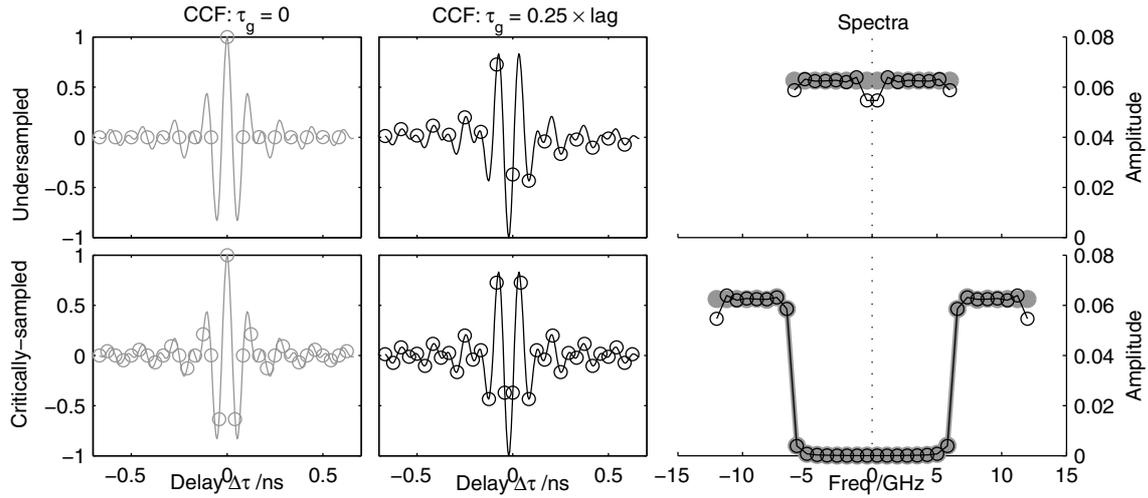
(a) The baseband signal (0–12 GHz) under and critically-sampled. The spectra show signs of aliasing. The cross-correlation function was sampled at 16 and 32 delays.
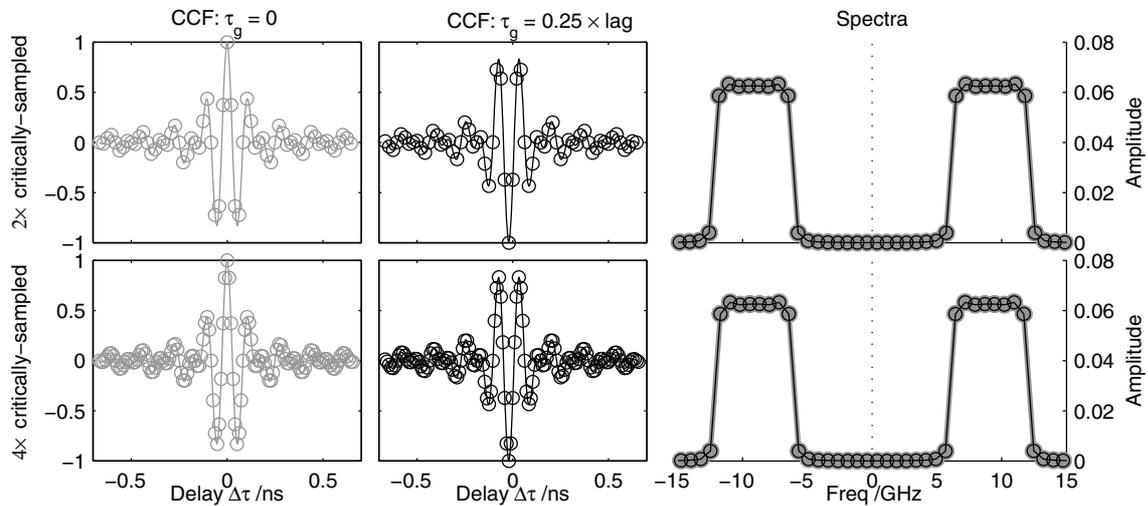
(b) Oversampling the baseband signal by 2× and 4× critical-sampling reduces aliasing.

Fig. 2. The effects of different oversampling rates on the recovered spectra. The left-most plots show the cross-correlation function from Eq. (7) as a function of the residual delay $\Delta\tau$. This is a snapshot when the source transits an east-west baseline so the geometrical delay $\tau_g = 0$ (case 1). The next set of plots to the right (case 2) show the cross-correlation function a little later when the centre of the envelope has drifted by a quarter of a delay spacing. The measurements at transit (case 1) are indicated by the grey circles and the measurements later on (case 2) are indicated by the black circles. The horizontal axis is the delay in $\Delta\tau$, which will be defined in Sect. 2.1. The right-most plots are the amplitudes calculated from the delay measurements. The filled grey marks are at transit (case 1) and the dark circles are a while later (case 2). The two sets of marks should match but it can be seen in **a)** that they do not. This mismatch shows up as the temporal modulation seen in Fig. 3. In *the upper row* of **a)**, although the 6 GHz signal bandwidth is critically-sampled at 16 delays, the 0–12 GHz baseband signal is undersampled. The two halves of the spectrum are side-by-side. The spectrum is aliased at both the high and low-frequency ends. In the lower row of **a)**, the highest frequency component (12 GHz) is critically-sampled with 32 delays. Now only the upper frequency range suffers from aliasing. In **b)**, the baseband signal is oversampled by a factor of 2 and 4 (sampled at 64 and 128 delays respectively). Now the two sets of spectral amplitudes match. Oversampling ensures a buffer between the two halves of the signal spectrum as well as with the image spectra. This reduces the spectral overlap and suppresses aliasing. In this case, oversampling the baseband frequency (0–12 GHz) by a factor of 2 is sufficient (an equivalent sampling frequency of 48 GHz at 64 delays). Oversampling above this does not reduce aliasing any further. As a minor technical detail, the forms of the spectra in **b)** are slightly different from the one in Fig. 1. Here (and also in Fig. 3), we shifted the spectra by half a sub-band using the shift theorem of Fourier transforms. This ensures that the sub-bands span the desired frequency range and is discussed further in Holler et al. (2007). Subsequent simulations will assume channel-shifting.

the phase of the local oscillator and it can be removed by phase calibration, at a later stage. τ_i is the total instrument delay difference in the IF inserted by the path compensator (τ_{pc}) and the differential delays inserted by the correlator lag (τ_{lg});

$$\tau_i = \tau_{pc} + \tau_{lg}. \tag{2}$$

The cross-correlation function (Eq. (1)) plotted against the correlator delays τ_{lg} (at ν_{IF}) consists of a cosine term bound by a sinc envelope (see Fig. 1a). The Fourier transform of the cross-correlation function gives the spectrum. So if we can sample the cross-correlation function sufficiently at regular correlator delays τ_{lg}, we can measure the spectrum of the source. This is the
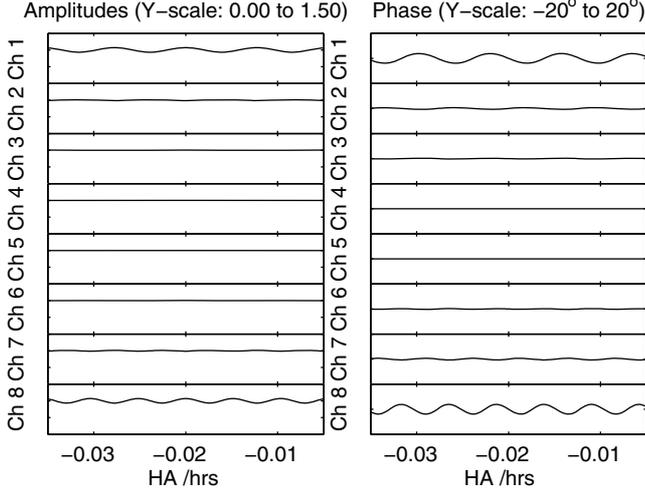
Fig. 3. Simulated time-stream data of the amplitude and phase for each spectral channel. The cycles are caused by aliasing and are worst in the edge frequency channels. This simulation is for a *real* correlator with parameters outlined in Table 1. Similar effects are also seen in the complex correlator.
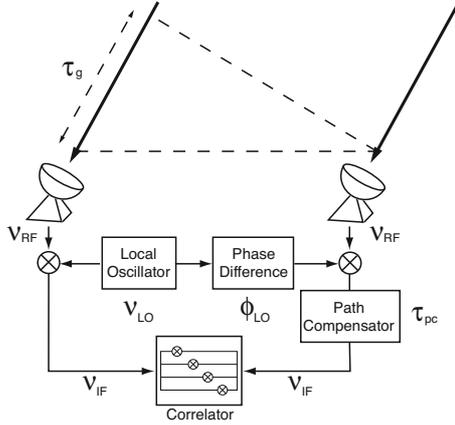

Fig. 4. A schematic of an interferometer system based on Thompson et al. (2001). The geometric delay (τ_{g}) between a pair of antennas is corrected by the path compensator delay (τ_{pc}) in the IF. The correlator samples the cross-correlation function at a number of delay steps (τ_{lg}).

basis of Fourier transform correlators that we will be discussing in this paper.

For brevity, we will define the residual delay at each correlator delay as

$$\Delta\tau = \tau_{\mathrm{g}} - \tau_{\mathrm{i}} = \tau_{\mathrm{g}} - (\tau_{\mathrm{pc}} + \tau_{\mathrm{lg}}). \tag{3}$$

We can then re-express Eq. (1) using $\Delta\tau$ and this will be the standard form we will use in the rest of the paper;

$$R_1 = R_0 \frac{\sin \pi \Delta\nu\Delta\tau}{\pi\Delta\nu\Delta\tau} \cos\left[2\pi\left(\nu_{\mathrm{LO}}\tau_{\mathrm{g}} \pm \nu_{\mathrm{IF}}\Delta\tau + \phi_{\mathrm{LO}}\right)\right]. \tag{4}$$

2.1. Real and complex correlators

Broadly speaking, there are two architectures for Fourier transform correlators. We will call these the real and complex correlators. The spectral information they give are the same but the way they sample the cross-correlation function is different. A more detailed discussion on Fourier transform correlator designs can be found, for example, in H07. The real correlator samples

at $2N$ delay steps and will produce N sub-bands. The complex correlator makes two measurements at each of the N delay steps; the in-phase component (like in the real correlator) and also the quadrature component. The quadrature component is detected by inserting a $-90°$ phase shift into one arm. For the quadrature component, the cosine in the cross-correlation function (Eq. (4)) is replaced with a sine;

$$R_2 = R_0 \frac{\sin \pi \Delta\nu\Delta\tau}{\pi\Delta\nu\Delta\tau} \sin\left[2\pi\left(\nu_{\mathrm{LO}}\tau_{\mathrm{g}} \pm \nu_{\mathrm{IF}}\Delta\tau + \phi_{\mathrm{LO}}\right)\right]. \tag{5}$$

We have also substituted in the residual delay $\Delta\tau$ defined in Eq. (3) for clarity. The complex correlator samples the in-phase and quadrature components at N delay steps. The two components can be conveniently expressed by a sum of the two orthogonal signals,

$$R_{\mathrm{x}} = R_1 + jR_2 \tag{6}$$

or written as an exponential;

$$R_{\mathrm{x}} = R_0 \frac{\sin \pi \Delta\nu\Delta\tau}{\pi\Delta\nu\Delta\tau} \exp\left[2\pi j\left(\nu_{\mathrm{LO}}\tau_{\mathrm{g}} \pm \nu_{\mathrm{IF}}\Delta\tau + \phi_{\mathrm{LO}}\right)\right]. \tag{7}$$

The delay steps of the Fourier transform correlator are determined by the total bandwidth of the correlator. For the critically-sampled real correlator, the delays are in steps of $T = 1/(2\Delta\nu)$. The complex correlator samples at only half-Nyquist rate ($T = 1/\Delta\nu$) because the cross-power spectrum of a complex cross-correlation function is single-sided (see Fig. 1b).

2.2. Fringe-rotation

Conventional fringe-rotation transforms the data so that the phase centre maps to the tracking centre. This is done by compensating for the cosine component of Eq. (4). Fringe-rotation can be applied either in hardware by controlling the phase of the local oscillator (ϕ_{LO}) or in software after recording the data. We will assume the latter case. The data must be sampled fast enough to avoid fringe-washing. For the complex correlator, the fringe data (Eq. (7)) is multiplied by

$$f_{\mathrm{fr}} = \exp\left[-2\pi j\left(\nu_{\mathrm{LO}}\tau_{\mathrm{g}} \pm \nu_{\mathrm{IF}}\Delta\tau\right)\right]. \tag{8}$$

This fringe-rotation factor can be calculated exactly from a combination of the geometry of the baseline (antenna pair), the position of the source and the frequencies ν_{LO} and ν_{IF}. Fringe-rotation will stop the phase-wrapping of the the fringes to give

$$R_{\mathrm{x}}f_{\mathrm{fr}} = R_0 \frac{\sin \pi \Delta\nu\Delta\tau}{\pi\Delta\nu\Delta\tau} \exp\left[2\pi j\phi_{\mathrm{LO}}\right]. \tag{9}$$

The Fourier transform of the stopped fringes gives the spectral data or visibilities. The visibilities can then be used to map the source.

For the real correlator (Eq. (4) and as illustrated by Fig. 1a), we cannot directly rotate the fringes in software because we do not have both the amplitude and phase information at each delay step. So the data must first be Fourier transformed across the delays into spectral channels (which is the usual operation for Fourier transform correlators). Each spectral channel is then fringe-rotated. The data from the complex correlator can also be fringe-rotated after the Fourier transform. This may be advantageous, for example when the fractional bandwidth ($\Delta\nu/\nu_{\mathrm{RF}}$) is high and the centre frequency used for fringe rotation is less constrained.

2.3. Spectral synthesis and mapping

In a conventional signal processing pipeline, the data recorded at each time sample are Fourier transformed across the delays to give the spectral data at each sub-band. For example, the continuous Fourier transform of Eq. (9) gives a rectangular spectrum as expected. The amplitude is calibrated against a flux calibrator source and corrected for changes in the system temperature (including effects like airmass). The phase is also calibrated against a phase calibrator source. This will remove phase offsets like ϕ_{LO}. In the ideal case, an unresolved source at the tracking centre will produce visibility data with a constant amplitude and zero phase. The successive visibility samples are time-averaged to reduce the data load and gridded to the two-dimensional aperture plane. This gives the spatial frequency components of the map. The coordinates of the visibility data in the aperture plane depends on the baseline geometry relative to the source and the RF of the sub-band. When all the visibilities have been gridded and appropriately weighted, the aperture plane is Fourier transformed to the map plane. The result is the dirty map – the convolution of the map with the synthesised beam. The synthesised beam of an interferometer usually has large sidelobes so the dirty map can be deconvolved using standard techniques such as CLEAN. Parts of the steps for reducing the data are summarised in Fig. 5.

3. Oversampling method

As we already showed in Sect. 1, it is possible to reduce aliasing by oversampling the cross-correlation function above the the Nyquist rate. This inserts a buffer zone in the spectral domain between the signal spectrum and the images. Using the new method that we will describe here, we could achieve this without increasing the number of detectors in the correlator.

The source is tracked as usual and the path differences between the antennas are compensated in discrete steps using path compensators. Normally, the data are fringe-rotated to align the phase centre with the tracking centre. This *stops* the fringes. Instead, we apply a modified fringe-rotation (Sect. 3.1) and allow the fringes to vary. When the path compensator steps, the data since the last step are collected together. The varying fringes trace out the cross-correlation function in Eq. (7) with $\tau_{\mathrm{g}} = 0$. The function is densely sampled so the Fourier transform (Sect. 3.2) introduces negligible aliasing[1].

3.1. Modified fringe-rotation

The aim of the method is to build up an oversampled set of data points for the cross-correlation function in $\Delta\tau$. But this is complicated because the exponential term in the response of the complex correlator (Eq. (7)) is also dependent on $2\pi\nu_{\mathrm{LO}}\tau_{\mathrm{g}}$. It is possible to remove just this term from the equation by applying a modified fringe-rotation function instead;

$$f_{\mathrm{mfr}} = \exp\left[-2\pi j\nu_{\mathrm{LO}}\tau_{\mathrm{g}}\right]. \tag{10}$$

The resulting fringe-rotated data will be a function of $\Delta\tau$;

$$R_{\mathrm{x}}f_{\mathrm{mfr}} = R_0 \frac{\sin\pi\Delta\nu\Delta\tau}{\pi\Delta\nu\Delta\tau} \exp\left[2\pi j\left(\pm\nu_{\mathrm{IF}}\Delta\tau + \phi_{\mathrm{LO}}\right)\right]. \tag{11}$$

---

[1] The oversampling method shares some analogies with drift scans. In a drift scan, the source is tracked but the path compensator is kept fixed. The spectrum can then be estimated by taking the subsequent Fourier transform of the time-series data.
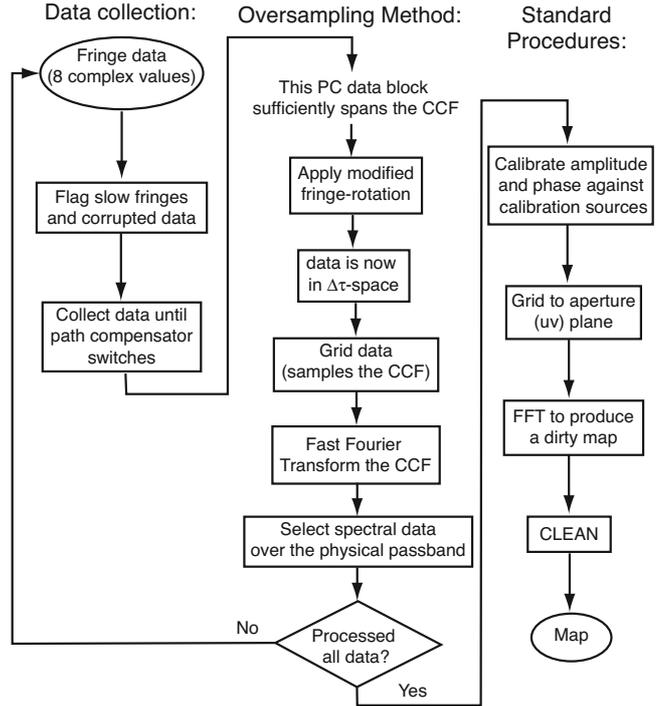

Fig. 5. A flowchart summarising the oversampling method.

By collecting consecutive time samples, we can build up an oversampled measurement of the cross-correlation function. We assume that the underlying signal is not changing over the time it takes to sample the whole cross-correlation function and we will return to this issue in Sect. 3.2.

In contrast, the real correlator is not very well suited to the oversampling method. The fringes cannot be directly fringe-rotated so it first needs to be Fourier transformed[2]. The spectral channels are fringe-rotated by the modified factor[3] f_{mfr} from Eq. (10) as before and then inverse Fourier transformed back;

$$R_1{}' = \mathrm{FT}^{-1}\left[f_{\mathrm{mfr}}\,\mathrm{FT}\left[R_1(\Delta\nu, \Delta\tau, \tau_{\mathrm{g}})\right]\right]. \tag{12}$$

Note that this gives the rotated fringes as a function of $\Delta\tau$. In addition to the computational burden, the Fourier transform step will cause aliasing. We will investigate this phenomenon through simulations in Sect. 4.2.1. This shortcoming for the real correlators may be overcome with a hardware fringe-rotation system. This approach also has the benefit of slowing down the fringes so that the data can be sampled at a lower rate. If fringe-rotation is to be applied in software, complex correlators are preferable because direct fringe-rotation is faster and much cleaner. However broadband $-90°$ phase shifters can be difficult to design.

So far we have assumed an ideal square passband. In reality, the source's spectrum may not be flat and the instrument may have a sloping passband. These factors may displace the centre-frequency and degrade the system's sensitivity. But following methods outlined by Thompson & D'Addario (1982)

---

[2] In signal processing, window functions are often applied to reduce sidelobes. But a window function should not be applied at this stage because the spectral resolution needs to be as narrow as possible to avoid aliasing – aliasing from the overlapping main lobes is a more acute problem.

[3] The conjugate relationship between the positive and negative spectra in a real correlator means that when applying fringe-rotation to a real correlator, the conjugate of Eq. (10) may have to be applied to one-half of the spectrum.

and Thompson et al. (2001), we calculated that the losses from fringe-rotation are negligible compared to losses from the sloping spectrum.

3.2. Spectral synthesis

We now have samples of the fringe-rotated cross-correlation function against $\Delta\tau$. We now divide the time-series data between path compensation changes so that in each *PC data block*, the cross-correlation function is completely sampled over the required range in $\Delta\tau$. The PC data block is then gridded and each pixel is weighted by the system temperature at the corresponding sample time. The gridded PC data block is then Fourier transformed[4] to give the oversampled spectrum. The wide spectral buffer between the signal and the images reduces aliasing. The synthesised spectrum over the passband is extracted, calibrated and gridded to the aperture (*uv*) plane as described before. This process is repeated for each PC data block. We have summarised the steps as a flowchart in Fig. 5 and we will describe some examples through simulations in the next section.

The period of the PC data block depends on the baseline and will vary over the course of an observation. A fundamental requirement is that the PC data block period needs to be short enough to avoid time-average smearing (see Thompson et al. 2001). So the periods of the PC data blocks T_{pc} must satisfy

$$T_{pc} < (d/D)/\Omega_E, \tag{13}$$

where d is the antenna diameter, D is the baseline and Ω_E is the angular speed of the Earth's rotation. The period of a PC data block for a complex correlator is given by:

$$T_{pc} = (\nu_{RF}/\Delta\nu)/\nu_f, \tag{14}$$

where ν_f is the fringe frequency. For real correlators, where the shortest path compensation bit is shorter, T_{pc} is half the above value. In AMI, when the fringe rate drops below $\nu_f \sim 10^{-3}$ Hz, it becomes difficult to separate the signal from slowly-varying non-astronomical noise, so these samples are usually rejected. Some additional data will have to be rejected at low fringe rates because they cannot completely fill the PC data block. But the proportion of lost data is small for a telescope like AMI.

4. Simulations

We tested the oversampling concept with simulations of a model telescope based on AMI. The observing frequency is 15 GHz with a 6 GHz bandwidth, mixed down with an LO of 24 GHz to give an IF band of 6–12 GHz. Path compensation is applied in the IF at discrete steps. The shortest delay line is the same as the correlator's delay steps so that the gaps in the cross-correlation function can be filled over the course of a single path step. Fringe-rotation is applied in software, rather than in hardware so the fringes need to be sampled sufficiently fast. For a relatively short 5 m east-west baseline, we expect fringe periods longer than 50 s, so sampling at 1 Hz is more than sufficient. 8 spectral channels are synthesised from the full 6 GHz bandwidth. We investigated both the real and complex correlators. The real correlator has 16 delay steps and the complex correlator makes a pair of measurements at each of its 8 delay steps. The parameters of the model telescope are summarised in Table 1.

---

[4] A window function could be applied to the gridded data before the Fourier transform step but it needs to be chosen with care; we discussed in H07 that window functions could systematically bias the estimated flux density of sources that are not at the tracking centre.

Table 1. Model telescope parameters.

RF	15 GHz (12–18 GHz)
Bandwidth	6 GHz
Synthesised sub-bands	0.75 GHz
LO (Lower sideband reception)	24 GHz
IF	6–12 GHz
Dish diameter	3.7 m
Baseline range	5–20 m
Simulated baseline	5 m east-west (250λ)
Integration	1 s
Spectral channels	8 (0.75 GHz each)
Latitude	52°
Test source declination	52°
Real correlator:	
Correlator delay steps	16 samples
Correlator delay steps	83 ps (25 mm)
Path compensator steps	83 ps (25 mm)
Tracking hour angle range	±56 s
Complex correlator:	
Correlator delay steps	8 complex samples
Correlator delay steps	167 ps (50 mm)
Path compensator steps	167 ps (50 mm)
Tracking hour angle range	±112 s

4.1. Simulations method

We simulate the fringes measured at each delay step using Eqs. (4) and (7). We modelled an unresolved source at the centre of the field but the method will work for any sources because we are simply measuring the cross-correlation function. The antennas track the point source about the transit (when the source is due south) over an hour angle range of ±56 s for the real correlator or ±112 s for the complex correlator. This fills the CCF function over a range of $\Delta\tau = \pm0.61$ ns. The time-stream data is then fringe-rotated as outlined in Sect. 3.1. The cross-correlation function is now a function of $\Delta\tau$ and we grid the data points to 256 pixels (see the upper set of plots in Fig. 6). We now have an oversampled measurement of the cross-correlation function. If the pixels are too narrow, some pixels may not contain any data points and this will degrade the spectrum. At the same time, the pixels must be fine enough so that the cross-correlation function is not washed out. Individual pixels could be weighted by factors such as the number of data points that contribute to the pixel and the system temperature associated with the data points. Finally, we apply the fast Fourier transform (FFT) to the gridded data and extract the signal spectrum over the 6–12 GHz passband (lower plots in Fig. 6). The amplitude and phase can be calibrated against an astronomical source by standard methods.

4.2. Results and discussions

4.2.1. Ideal correlator

The recovered spectrum is in steps of 0.75 GHz and is the same as when the conventional *one-shot* method is used. The oversampled spectrum extends from $\nu_{IF} = $ DC to 104 GHz and only a portion is shown in Fig. 6. The spectrum for the complex correlator in Fig. 6a does not appear to suffer from aliasing. The arguments in Sect. 1, together with Fig. 2 suggests that the oversampling method should reduce aliasing.

In the real correlator, fringe-rotation introduces aliasing via the Fourier transform. This can be seen as small cycles in the data points (Fig. 6 and details in Fig. 7). These arise from the time-domain alias cycles in Fig. 3. Fringe-rotation aliasing is
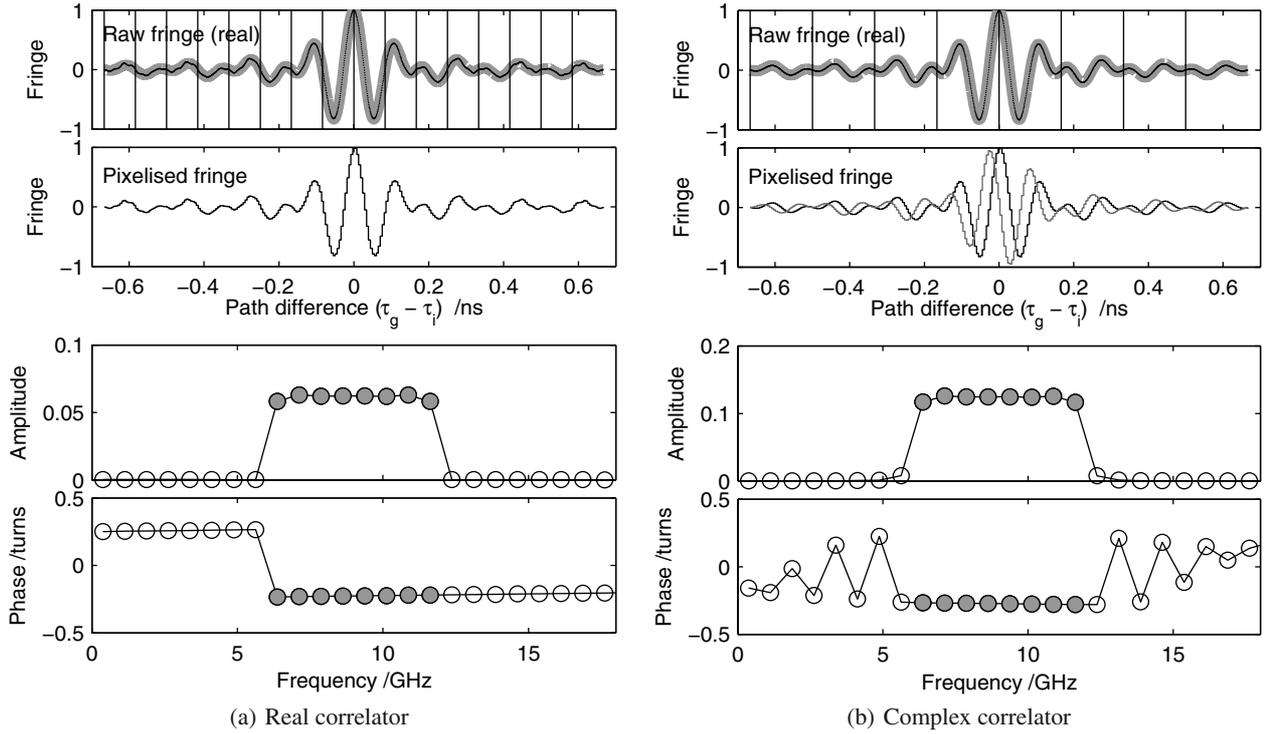
Fig. 6. The oversampled spectral synthesis for a real correlator **a)** and complex correlator **b)**. The top plots are the real components of the fringes. (1) The thick grey lines mark the ideal cross-correlation function we would expect from Eq. (4) for **a)** and the real component of Eq. (7) for **b)**. The geometrical delay τ_g for the thick grey curves are set to 0. The dark dots (seen as continuous curves at this resolution) are the modified-fringe-rotated data points measured by the correlator at each time sample (from Eqs. (12) and (11) respectively). In **b)**, the imaginary components have not been shown for clarity. The vertical lines indicate the range in $\Delta\tau$ spanned by each delay; the curve in each section is traced out from the left boundary to the right. (2) In the second plot, the modified-fringe-rotated data points are pixelised into 256 pixels. For the complex correlator, the dark curve is the real component of Eq. (11) and the grey curve is the imaginary component. (3) The pixelised data are fast Fourier transformed. *The bottom two plots* show the amplitude and phase of the spectrum. The estimated spectrum extends to higher frequencies so only a portion is shown. The filled circles mark the synthesised spectral channels of interest over the physical passband of the system.

worse near the edge of the cross-correlation function. The root mean square (rms) deviation of the recovered spectrum from the ideal case is negligible (below 0.1 percent). But if there are delay errors, the rms deviation may become significant and this will be discussed in Sect. 4.2.2. If using the oversampling method on real correlators, we would recommend hardware fringe-rotation. If software fringe-rotation is chosen, as shown here, the effects of fringe-rotation aliasing needs to be investigated further. Gridding is equivalent to down-sampling and the conventional practice of introducing a lowpass anti-aliasing filter before gridding may be beneficial. With larger fractional bandwidths ($\Delta\nu/\nu_{RF}$), the small cycles from fringe-rotation aliasing will become longer relative to the pixel size.

4.2.2. Correlators with delay errors

The oversampling method also corrects for delay errors prevalent in analogue Fourier transform correlators. It is possible to determine the delay errors either by bench measurements or by tracking a source with the path compensator held fixed (see Harris & Zmuidzinas 2001; H07). The delay errors are incorporated into calculating $\Delta\tau$ and gridded uniformly so the recovered frequency channels are also uniformly spaced.

Figures 8a and b show simulations for both the real and complex correlators with delay errors of up to 10 percent of the delay steps. For the complex correlator, we assumed that the delay errors for each pair of detectors are the same. This is a reasonable

assumption for our correlator architecture because we found that most delay errors arise in the long delay lines common to both the real and imaginary signals (see H07). In both correlators, there are gaps in the samples from the irregular delay steps. The recovered spectrum is given by the spectral response of this distorted window function convolved with the true spectrum. This will inevitably worsen the inter-channel spectral leakage.

The presence of delay errors makes aliasing resulting from fringe-rotation of the real correlator data noticeably worse (Fig. 8a and details in Fig. 7). This is consistent with our previous findings (H07). When fringe-rotating the spectral channels from the real correlator, we assumed regularly-spaced centre-frequencies at the design frequencies. These may not be the best choice for an irregularly-spaced correlator and may even introduce small phase errors during fringe-rotation. Fringe-rotation aliasing gives rise to a low-level spread-spectrum noise in the amplitude plot (Fig. 8a). Perhaps the most significant evidence against applying software fringe-rotation to a real correlator is the potential for systematic errors. We found about 3 percent rms deviation from the ideal spectrum when there are delay errors. The alias cycles seem to be largely responsible for this because the rms deviation for a complex correlator with lag errors was less than 0.01 percent.

It is worth noting that a correlator with delay errors will fundamentally reduce sensitivity and no amount of signal processing can compensate for it. This is because the SNR of the recovered spectrum is optimised when sampled at Nyquist steps (or at
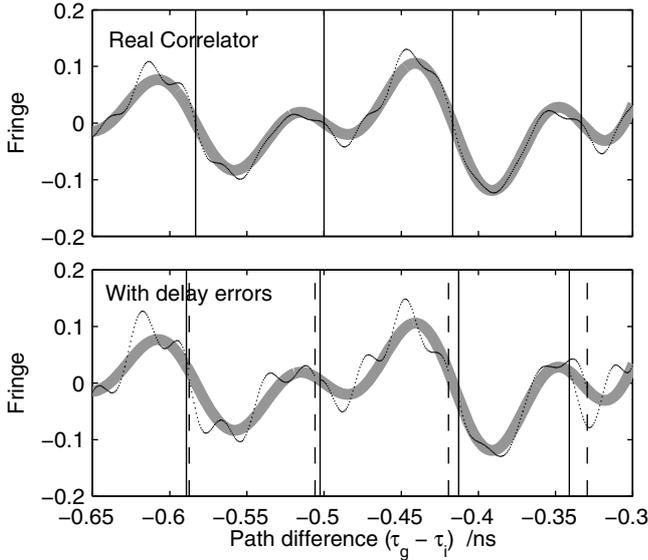
Fig. 7. Details of the reconstructed cross-correlation function for the real correlator in Fig. 6a towards the edge of the delay range. The thick grey curve is the expected function from Eq. (4) with τ_g set to 0. The dots are the data points generated from Eq. (12) after the modified fringe rotation. *The upper plot* is for a perfect real correlator (from Fig. 6a) and *the lower plot* is with delay errors (from Fig. 8a). The solid vertical lines indicate the position of each delay in $\Delta\tau$ at the beginning of the PC data block. The dashed vertical lines are the position at the end of the PC data block. Delay errors lead to missing samples or overlapping samples inside the narrow gaps indicated by these lines. Fringe-rotation aliasing is apparent in both cases and it is also worse when there are delay errors.

half-Nyquist steps for complex correlators). Delay errors cause the noise components between the delay steps to be correlated, so the measurements are no longer independent and the SNR is degraded. The benefits of applying the oversampling method in the presence of delay errors are: (1) the recovered spectral channels are regularly gridded and are at the designed centre frequencies; (2) normally, delay errors worsen the effects of aliasing (H07). Figure 8b and the rms figure for the complex correlator suggest that the oversampling method reduces aliasing, even in the presence of delay errors.

4.2.3. In the presence of noise

In Figs. 8c and d, we added independent Gaussian random noise to the fringes of a complex correlator. Noise in each time-frame are independent and it can be shown that in our idealised system, the noise between delay steps are also independent. The recovered spectra are noisier as expected but the oversampling method is robust in the presence of noise. It is also a linear process so successive data sets can be stacked (Fig. 8d). Because the noise is independent between samples and also between the delay steps, the noise energy is spread over the whole of the oversampled spectrum. We only select spectral channels over the passband and discard the rest. This does not improve the SNR beyond what can be achieved by standard methods. The SNR improvement from applying this matched filter is the same as from averaging successive samples in the standard methods.

4.3. Issues in a practical system

Several issues will need to be addressed when using the over-sampling method in a practical system and design tolerances must be set. Firstly, more detailed simulations of the method

should be conducted to quantify the reduction in aliasing. So far, we have assumed that the passbands of the detectors at each delay steps are identical. Any variation will restrict the effective bandwidth and degrade the sensitivity. If there are variations in the passbands, each detector will trace out a different cross-correlation function because the cross-correlation function is related to the Fourier transform of the passband. When the data points are merged to reconstruct the full cross-correlation function, there will be discontinuities between the detector outputs. However the degradation is probably comparable to standard methods.

If the path compensators do not step exactly in multiples of one delay step (due to manufacturing tolerances in the path compensators), there will be unsampled gaps in $\Delta\tau$. Optimising the path compensation for an array can be more complicated and some delay ranges ($\Delta\tau$) may not be fully sampled. Alternatively, some samples may have to be rejected because they are corrupted. These missing samples will degrade the overall sensitivity in the same way as delay errors. Making the grid coarser may mitigate some of this but at the expense of washing out the cross-correlation. The data needs to be stable over the PC data block period. This could be worked round by opting for real correlators (with hardware fringe-rotation), which cuts down T_{pc} by a half.

4.4. Application to AMI

Although both the real and complex correlators were trialled for AMI, the real correlator had been chosen (see H07). As demonstrated here, the oversampling method is not suited for real correlators relying on software fringe-rotation. Later on, it was found that the bottom two frequency channels in AMI (12−13.5 GHz) suffered from geostationary satellite interference. The effects became noticeable when observing at low declinations and it was decided to block this frequency range all together with hardware filters. Although this reduces the bandwidth of the instrument, it also goes some way towards reducing the aliasing on one side of the spectrum. This change avoided the need for the more complicated oversampling method and subsequently, we did not implement it in the AMI data reduction pipeline.

5. Conclusions

We have described a software-based method for overcoming aliasing in critically-sampled Fourier transform correlators. The method reduces aliasing by reconstructing the oversampled cross-correlation function from successive samples. Data from complex correlators can be fringe-rotated directly and efficiently in software. If using this method on real correlators, the signal should be fringe-rotated in hardware to avoid systematic errors.

The edge spectral channels are usually discarded because they are the worst affected by aliasing. For an 8-channel complex correlator simulated here, the sensitivity could be improved by up to 15 percent by retaining the edge channels. The additional computational overhead and complexity are offset by an improvement in sensitivity and reduction in systematic errors from aliasing. This could be a significant benefit for Fourier transform correlators with less than 10 spectral channels. Some data with low fringe rates will have to be rejected, but for AMI-like telescopes with fractional bandwidths $\Delta\nu/\nu_{RF} \gtrsim 0.4$ and maximum baseline to dish ratios of $D/d \lesssim 10$, the fraction of data lost is relatively small. We have also shown that we can naturally compensate for delay errors to recover regularly-gridded frequency channels at the design frequencies.
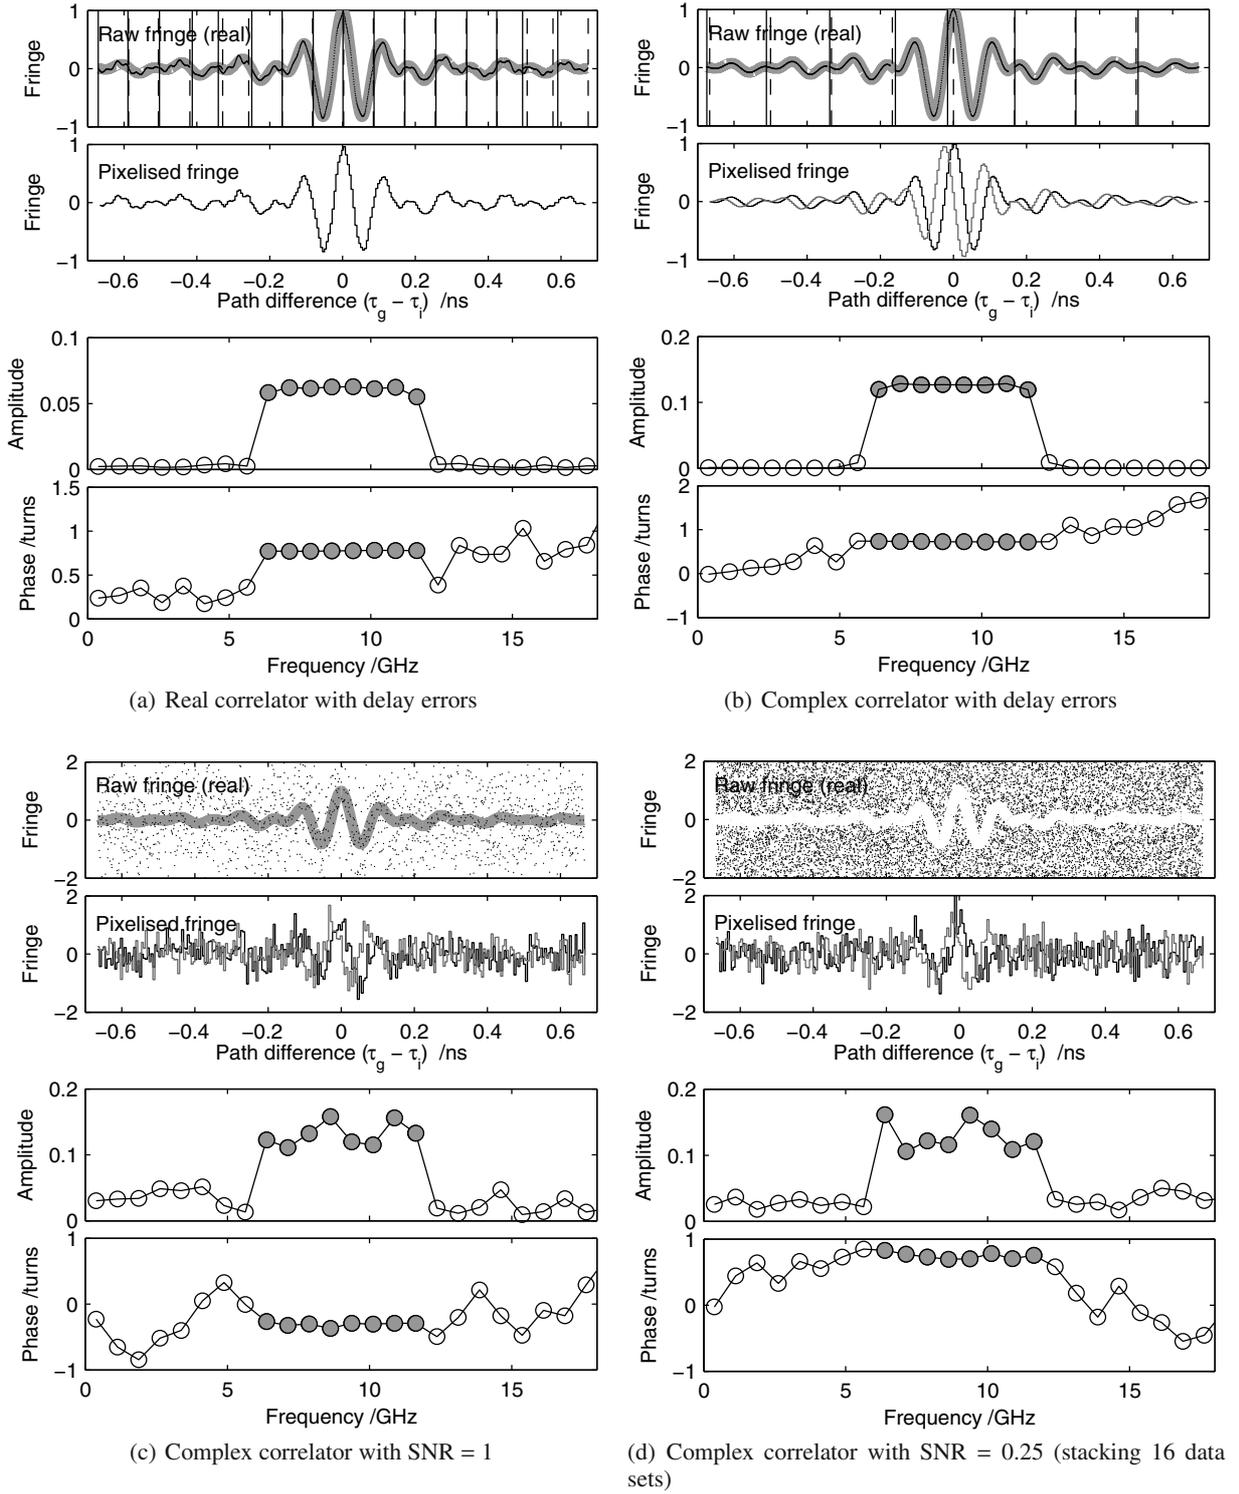
(a) Real correlator with delay errors

(b) Complex correlator with delay errors

(c) Complex correlator with SNR = 1

(d) Complex correlator with SNR = 0.25 (stacking 16 data sets)

Fig. 8. a) and **b)**: The oversampling method in the presence of delay errors. In both cases, the errors are up to 10 percent of the delay steps and there are visible gaps in the data. The solid vertical lines mark the left-most range of each delay and the dashed vertical lines mark the right-most range. The delay errors can be clearly seen from the relative alignment of these lines. The other effect in the real correlator **a)** is the higher level of fringe-rotation aliasing in the cross-correlation function (see details in Fig. 7). The thick grey curves are the ideal cross-correlation curves from Eq. (4) with τ_g set to 0. The dark points are the modified-fringe-rotated data points generated from Eqs. (12) and (11). **c)** and **d)**: Gaussian random noise was added to the fringes of a complex correlator. The instantaneous SNR is quantified by the ratio of the peak of the cross-correlation function to the standard deviation of the noise. In **d)**, 16 data sets are stacked so that the recovered spectrum has the same SNR as **c)**. The phases have been unwrapped arbitrarily.

References

Harris, A. I., & Zmuidzinas, J. 2001, Rev. Sci. Inst., 72, 1531

Holler, C. M., Kaneko, T., Jones, M. E., Grainge, K., & Scott, P. 2007, A&A, 464, 795

Kaneko, T. 2006, in Proc. SPIE Ground-based and Airborne Telescopes, ed. L. M. Stepp, 6267

Li, C.-T., Kubo, D., Han, C.-C., et al. 2004, in Proc. SPIE 5498, ed. C. M. Bradford, P. A. R. Ade, J. E. Aguirre, et al., 455

Roberts, P. P., Leach, M. R., & Wilson, W. E. 2007, PASP, submitted

Sunyaev, R. A., & Zel'dovich, Y. B. 1972, Commun. Astrophys. Space Phys., 4

Thompson, A. R., & D'Addario, L. R. 1982, Radio Sci., 17, 357

Thompson, A. R., Moran, J. M., & Swenson, G. W. 2001, Interferometry and Synthesis in Radio Astronomy, 2nd edn. (John Wiley & Sons)

Zwart, J. T. L., Barker, R., Biddulph, P., et al. 2008, MNRAS, submitted, [arXiv:0807.2469]