ALMA Memo #305
A digital BBC for the ALMA interferometer

G. Comoretto
Osservatorio di Arcetri
Florence, Italy

April 26, 2000

Abstract

A digital baseband converter is proposed for the ALMA correlator. The digital local oscillator/mixer can be implemented with lookup tables, and increases only slightly the complexity of the digital filter board.

Sideband rejection can be implemented either in the filter section, at the expense of a roughly doubled number of filter taps, or adopting a complex correlator. Several approaches are possible, that impact to a various extent the existing correlator architecture. In the simplest approach, only the memory board needs to be modified.

The digital nature of the mixer causes a loss in signal to noise ratio of a few percent. More important is the presence of LO harmonics, that create ghost images of the input lines. Spur level depends mainly on the number of bits in the data path between the mixer and the filters. For a 3-bit mixer, a spur level around -31 dBc (decibel with respect to the fundamental) can be obtained. To achieve a -38 dBc spur level, a 4 bit mixer is required.

1 Introduction

In MMA memo 204, a digital filter has been proposed to replace all BBC filters in the ALMA interferometer [6]. In this approach, the signal is first sampled at 4 GS/s (Gigasamples per second), parallelized in 32 streams at 125 MS/s each, digitally filtered and resampled.

Inserting a multiplier in the digital stream, before the filter, would allow for digital downconversion. This would completely eliminate the need for a baseband converter (BBC) before the sampler, and would simplify the filter unit.

It would however require a complex correlator, while the ALMA interim correlator is intrinsically real. It is possible to overcome this limitation with minor modifications in the software and in the memory buffer board, but this is not possible for the maximum and minimum bandwidth of the original design.

In this report, the proposed digital BBC is analyzed first in the overall architecture, and then in detail (chapters 2 and 3). A summary of advantages and disadvantages of the various options is presented in chapter 4. Effects due to the repeated quantization and to the digital nature of the local oscillator are examined and discussed (chapters 5 and 6). Preliminal simulation results are presented in chapter 7.

The design for the ALMA correlator and baseband converter is currently in an advanced stage, and such radical changes in the overall architecture as eliminating the local oscillators probably cannot be seriously considered. The considerations made in this report are however quite general, and could be useful, for example in designing specific elements of the European advanced ALMA correlator.

2 Design principle

The BBC is required before the correlator for several reasons:

- band selection: several BBC's can observe independent spectral features in the same 2 GHz broad band signal.
- fringe rotation: the BBC oscillator can be programmed to compensate for Doppler shift of one antenna with respect to the reference point of the array.
- SSB cancellation schemes: by offsetting by an adequate value both the front end and BBC oscillators, one can in principle suppress the undesired sideband of the front end mixer (see MMA report 190 [5]).

Several possibilities to implement the BBC, fully or partially, with digital techniques, have been proposed.

2.1 Digital filters in the current design

In the proposed ALMA backend, the BBC is composed of a DDS-controlled oscillator, and a conventional SSB converter.

The signal is then sampled and quantized, and processed by a finite impulse response (FIR) filter[6], with length proportional to the inverse fractional bandwidth. A length of 96 taps has been shown sufficient to obtain a useful bandwidth of 95% for a 1 GHz filter, and this length is increased by a factor of 2 for each factor of 2 reduction in the output bandwidth [8].

Since the correlator operates in parallel over multiple 125 MS/s data streams, a total of 16 filters is required for a 1 GHz bandwidth, for a total of 1536 taps. The number of taps is constant for reduced bandwidth, since parallel filters are cascaded together.

2.2 Bandpass digital filter

The analog BBC and sampler have poor performance at near-zero frequencies. It is therefore advisable, for reduced bandwidth, to select a band different from the baseband. The band selection filters then operate in one (selectable) frequency slot that, after resampling at the correct final frequency, is aliased to the baseband. If the final sampling frequency is f_s, any band from $kf_s/2$ to $(k+1)f_s/2$ can be selected, with k integer.

A bandpass filter is not much different, in hardware, from a low-pass one. The main difference is that the required number of taps must be doubled, because the transition region, not usable for observations, occur now both at the low and the high end of the bandwidth.

Therefore one must provide a total of 3072 taps, for a 95% usable band, or degrade the latter to a 90%.

2.3 Digital LO with digital phase shift

It is possible to closely simulate the analog BBC with a digital system. The design principle is shown in fig. 1.

The signal is downconverted to a 2 GHz band by a fixed LO, and sampled. The sampled signal is beated with a digital representation of a sinusoid, computed in a lookup table from phase information generated by a Direct Digital Synthesizer (DDS). Two signals in quadrature are generated, and the two corresponding signals (real and imaginary) are separately filtered, differentially phased by 90 degrees, and then added together. The differential phase shift of 90 degrees is generated in the band pass filters.

It is not difficult to introduce a prescribed phase shift in a bandpass filter, if the upper and lower bounds are not too different. Probably it will not be necessary to increase further the total number of taps to accomodate this, but some simulation work is needed. However, separate filtering for the real and imaginary paths is needed, doubling the total number of taps to approximately 6000.

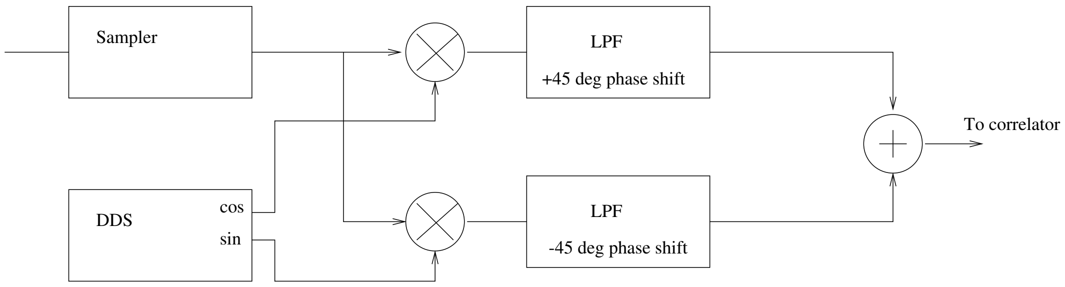

Figure 1: Implementation of a digital BBC using digital phase shift networks.

This approach has been proposed by other groups, and will not be considered further in this report. However, most of the considerations on the effects of a digital LO, described in chapter 5 and 6, are relevant also for the BBC with digital phase shift described here.

2.4 Digital LO with complex correlation

A digital quadrature converted signal, as obtained in the previous design, can be analyzed by a complex correlator. As in this latter, the 2 GHz bandwidth, sampled at 4 MS/s, is multiplied with a quadrature digital LO (fig. 2).

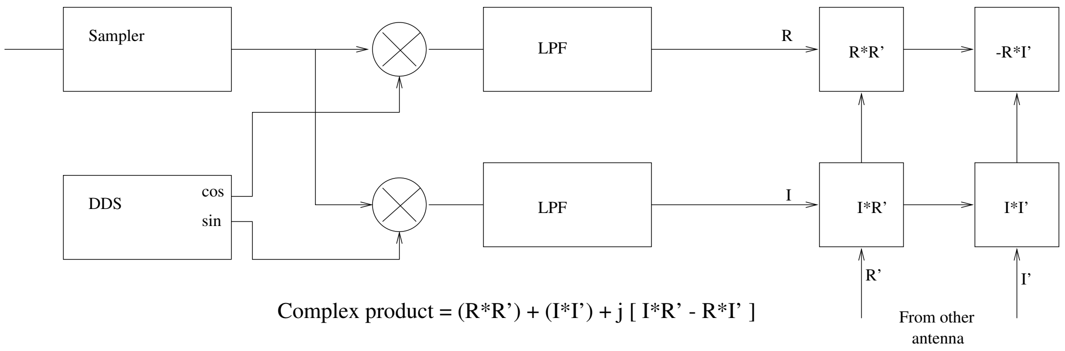

Figure 2: Digital BBC conceptual scheme. The four multipliers represent four sets of correlation planes

The two quadrature data streams, again separately filtered at a frequency $f_s/2$, carry information relative to a total bandwidth f_s (from $-f_s/2$ to $+f_s/2$. Sideband separation (positive and negative frequencies) is done in the correlator, and is limited only by numerical inaccuracies, well below the noise floor. This approach requires filters with a cutoff frequency at half the required bandwidth, and therefore with a doubled number of taps with respect to a low pass filter in a conventional correlator. However, only low-pass filters (not band-pass) are needed, and then the total number of taps is equal to that of a corresponding bandpass filter with the same usable bandwidth (approximately 3072 total taps).

A complex correlator is then required to analyze the signal. Complex correlation requires a double number of channels, but processes a double bandwidth with respect to a real time multiplexed correlator at the same clock frequency. Therefore both the total number of FIR taps and the number of (real) correlation channels are exactly the same than for the bandpass, analog BBC, design.

A complex correlator is very similar to a multiplexed one. Real and imaginary delays are computed in separate multiplication chains, that are added in the postprocessing. However the ALMA interim correlator [4] is not time multiplexed. Signal samples are analyzed in their natural time sequence, although at reduced speed. It is still possible to compute a complex correlation, if there are at least 4 separate correlation planes available for each frequency (time multiplexing) slot. In the current ALMA design, this means a usable bandwidth of 1 GHz or less (real bandwidth of 0.5 MHz). In this case, the digital BBC produces two streams of data at 1 GS/s, and four sets of correlation planes are used to analyze each of the four possible products RR', $-RI'$, IR', II' (fig. 4).

One of these products (-R*I') must have its sign changed. This may be done either in the correlation chip, or in the readout logic. Post processing software must also be modified to deal with complex correlation functions.

Alternatively, the correlator unit cell may be modified in order to deal with complex correlations. This does not appear feasible for the interim correlator, but could be considered for the advanced design (see section 3.2.2.

3 Implementation

The main additional elements for a digital BBC are the LO/mixer and the complex correlator. Those elements are described below. The LO/mixer can be easily implemented using RAM lookup tables in a programmable logic. Some different schemes to implement a complex correlator using existing hardware are proposed.

3.1 Digital LO and multipliers

Figure 3 shows a block diagram for the proposed implementation of a BBC.

A DDS register provides a LO phase with a rate of 125 MHz. The phase must be sufficiently accurate to avoid that the reconstructed waveform does not degrade the signal accuracy, and does not generate spurious harmonics. For a 3-level signal quantization, at least 6 bit phase values are required. In chapter 6 it will be shown that spur level depends on the number of phase bins. To obtain a better spur level, 7 bit phase values are used. A separate phase must be computed for each time multiplexed sample from the digitizer. This phase may be computed from a common 12-bit phase value by a separate adder, that adds the correct phase offset for the current frequency corresponding to the appropriate number of 4 GHz sampling periods.

The mixer is implemented as an array of RAMs, that have as input the sampled signal, together with the phase signal from the DDS. The mixer operates separately on each parallelized data stream, at the reduced speed of 125 MHz. The RAMs contain precomputed lookup tables to implement at the same time both a SIN/COS table, a multiplier and a resampling unit for the converted signal (fig. 3).

A total of 64 lookup tables implement the functions $C = A\sin(P + Np)$, and $C = A\cos(P + Np)$, where A is the input signal, P is the DDS phase, p is the phase increment corresponding to 1/(4GHz), and N is an integer from 0 to 31 specifying the time multiplexed channel. Each lookup table is a 1024×3 RAM, computing a 3-bit function from a 3-bit signal and a 7-bit phase. The RAM content must be updated every time the frequency of the DDS is changed. To allow frequency changes on the fly, a dual-bank system is used. This can be easily implemented using an extra address bit.

Memory usage can be reduced representing input data as sign-magnitude. The result sign can be computed with a exclusive-or logic, and the RAM LUT table dimension is reduced to 256×2 (2-bit signal, 6-bit phase).

Each RAM output replaces the corresponding input. At the output of the mixer, one has two streams of 32×3 bits each. Each stream feeds 8 digital low-pass filters, that reduce the signal bandwidth by a power of 2. The maximum available bandwidth is 0.5 GHz, corresponding to an effective analyzed bandwidth of 1 GHz. To use the full 2 GHz band, the digital LO and filter board is effectively bypassed, e.g. by placing unitary coefficients in the appropriate registers. The minimum bandwidth is 62.5 MHz (complex), corresponding

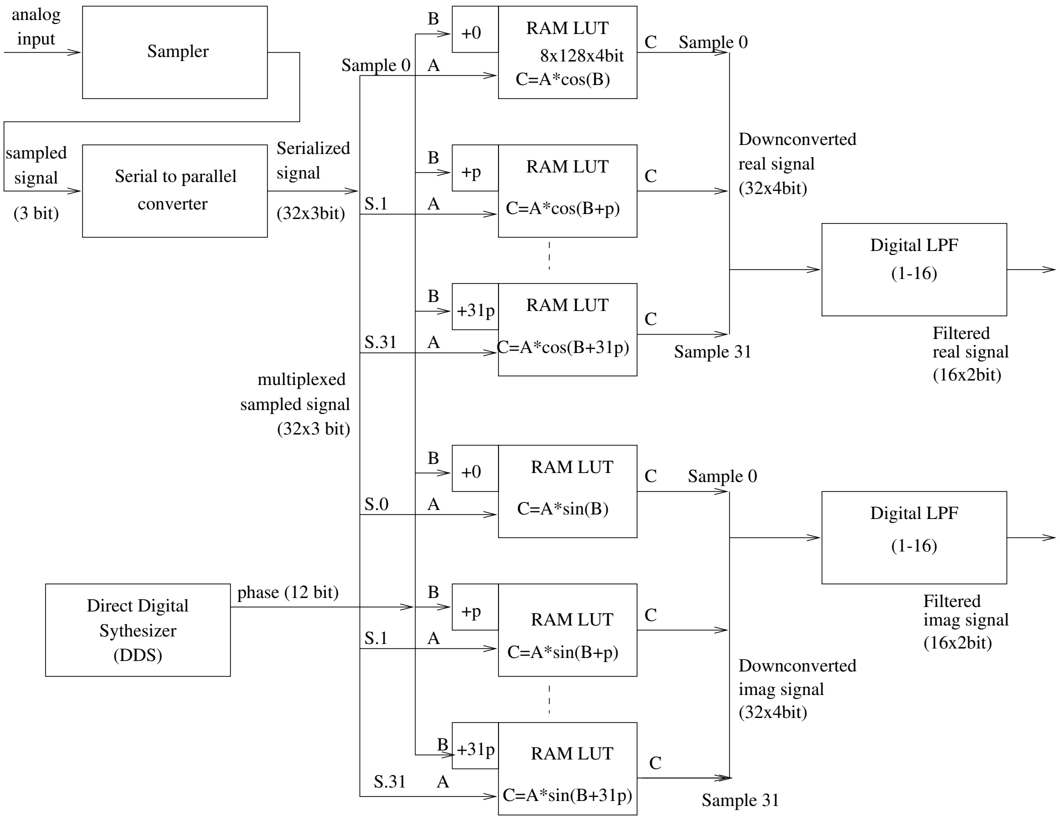

Figure 3: Possible implementation of a digital BBC. The sampler and filter sections are almost identical to the original digital filter concept.

to a 125 MHz input signal. Low-pass filters can be cascaded to produce an increased length FIR filter at reduced bandwidth, as in the original design.

Several Field Programmable Gate Array (FPGA) chips include memory in the chip, with almost the appropriate granularity. A full bandwidth converter can be implemented, for example, using a couple of XCV1000 Xilinx chips. The same chips can also implement the DDS oscillator.

3.2 Complex correlator

A complex correlator must compute the quantity $(RR'+II')+j(IR'-RI')$, where $j=\sqrt{-1}$, R, I are the real and imaginary data streams from the first antenna, and R', I' those relative to the second one (see figure 2).

These four correlation products can be computed in four separate real correlators. This is not very different from the case of a time multiplexed correlator, in which correlation products are computed between interleaved samples of the data stream. Considering that a complex correlator processes a doubled band with respect to a real one, and the number of correlators increase with the square of the frequency, a real and a complex correlator for the same bandwidth requires the same number of individual correlators.

3.2.1 Complex correlation in ALMA baseline correlator

However the ALMA correlator does not process an increased bandwidth in this way, but slows down the original data stream, and divides it into small contiguous frames, that are then processed in parallel. A total of 32 correlator planes are required to process the full 2 GHz bandwidth. If the bandwidth is reduced, different planes process different delay ranges, for increased frequency resolution.

In this case is still possible to compute the four products required for a complex correlation if the total required bandwidth is 1 GHz or less, using separate correlator planes (fig. 4).

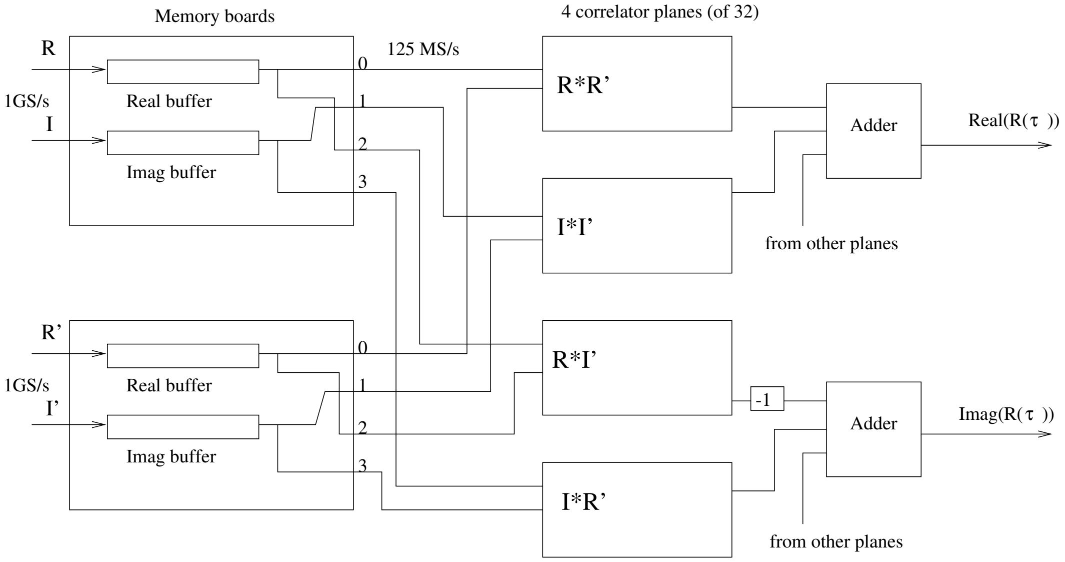

Figure 4: Implementation of a complex correlation in the ALMA proposed interim correlator. This is possible only if at least two independent set of planes are available for each time slot, i.e. for a total bandwidth of 1 GHz or less. Modifications need to be implemented in the memory recirculation board only.

For example, a 1 GHz bandwidth is represented by a 500 MHz complex data stream, and each product can be processed by a 500 MHz correlator (composed of 8 planes). The total number of delay channels, and frequency resolution, is the same of a real correlator at 1 GHz bandwidth.

Correlation products for each plane are added by ad hoc electronics. The same electronic can add together the real and imaginary parts of the complex correlation product. The two imaginary products have opposite sign, and the readout logic must be modified in order to complement the products from the relevant planes.

The memory board must also be modified, in order to separate the real ad imaginary streams from the sampler/LO/filter boards. Both these modifications appear not difficult to implement in the proposed correlator architecture.

3.2.2 Complex correlator chip

A more radical approach would be to modify the correlator chip, in order to make it capable of complex correlations. This does not appear feasible for the current, already designed, chip, but could be an option for an advanced correlator.

In this approach, the data stream consists of alternate real and imaginary samples. Therefore, the memory board operates exactly as for a real correlator, and no limitations on the bandwidth are needed.

The chip is very similar to a real correlator, i.e. is composed of a shiftregister, that delays the signal by one sample for each clock cycle, and a set of multipliers/adders, that compute and integrate the correlation function (fig. 5).

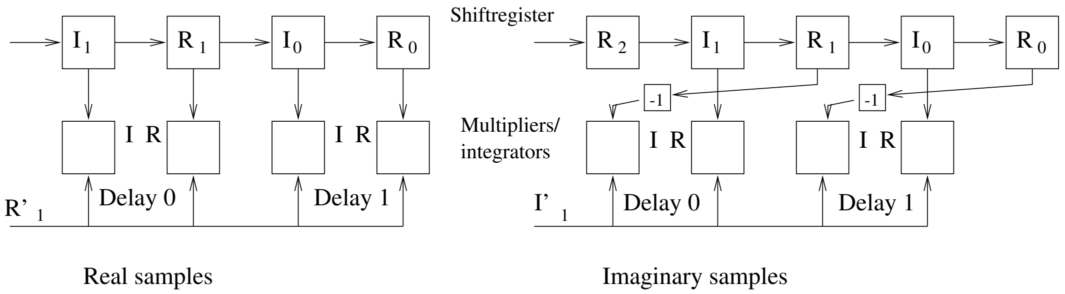

Figure 5: Architecture of a complex correlator. The data stream is composed of alternate real and imaginary samples. The undelayed input is delayed by one (real+imaginary) sample. The accumulators compute alternatively the real and imaginary part for each correlation delay. The correlator computes products RR' and IR' in the first cycle (left), and products RI' and II' in the second (right)

Since each data sample is composed of two elements, two clock cycles are required for each sample. The chip, then, processes the data at half speed, but this is compensated by the doubled band of a complex correlator.

In the first cycle, the real sample R' is presented at the undelayed (prompt) input, and the terms RR' and IR' are computed. In the second cycle, the imaginary sample I' is processed, giving the terms $-RI'$ and II'. The latter is just added to the real accumulators. To compute the former, the imaginary multipliers must have a multiplexer at the delayed input, that recovers the R sample now delayed by two shiftregister units. It must also complement its value, as required.

The prompt input must be delayed by one (complex) sample, requiring an extra shiftregister cell before the prompt input distribution. An extra cell must be added at the end of the shiftregister to hold the R sample during the second multiplication cycle (see fig. 5).

This correlator behaves for any other aspect like a conventional real correlator. Units can be cascaded, paralleled, etc. in the same way, each data stream carries the same amount of information, and the total bandwidth and frequency resolution depends in the same way on the number of shiftregister, multiplier and adder units employed.

4 Advantages and disadvantages

A digital BBC presents several advantages over an analog counterpart. Digital systems are in general more reliable and stable over their analog counterparts. The overall system complexity added is very small, a fraction of that introduced by the digital filter concept, and replaces a significant amount of analog hardware.

The critical 90° degrees phasing network required for a SSB is no more required. The fully digital LO is more stable and less critical with respect to a DDS followed by some analog system (DAC converter, frequency multiplier).

A disadvantage of the proposed system is that the maximum bandwidth must be dealt separately, without any conversion. This imposes the presence of a DDS before the sampler, at least for fringe rotation and phase switching, and therefore most of the advantages quoted above are reduced or lost.

Also the *minimum* bandwidth available is doubled with respect to that of an analog BBC. A complex correlation always require four separate correlators. If for higher bandwidth this does not introduce penalties, below 125 MHz bandwidth one stops trading channels for bandwidth, and the maximum number of spectral channels is reduced by a factor of 2. A more complex recirculation scheme may be employed for keeping the product $B{\cdot}N_{ch}$ constant, but this is a significant modification of the correlator design.

Both these limitations do not occur for a full digital BBC (including phase shift network), and for a complex correlation chip.

Quantization of the DDS sinusoid and phase introduces spurious conversion products. For a 3-bit quantization, LO spurs and spurious conversions appear at -30 dBc approximately. The two effects are independent: LO spurs occur even with only noise at the input, while spurious conversions are replica of strong signals, translated in frequency and attenuated. Both effects are reduced, but slowly (6dB/bit), increasing sinusoid quantization bits. This problem is dealt more extensively in chapter 6.

In this context it should be noted that any DC component in the quantized signal (e.g. due to a small offset in the sampler thresholds) is seen as a strong line at zero frequency, and produces undesired components in the output spectrum.

5 Quantization efficiency

In this approach, the signal is quantized several times: at the sampler, after the frequency conversion, and after the final FIR filter. Any quantization scheme introduces a loss in the SNR of the spectral extimate, and must be carefully considered.

A detailed computation of this loss is given by Cooper. He computed the correlation function $R(\rho)$ obtained from a digital representation X_1, X_2 of two analog signals x_1, x_2, distributed with a binormal distribution of correlation coefficient ρ.

$$P(x_1, x_2; \rho) = \frac{1}{2\pi\sqrt{1-\rho^2}} \exp\left(-\frac{x_1^2 + x_2^2 - 2\rho x_1 x_2}{2(1-\rho^2)}\right) \tag{1}$$

The complete determination of $R(\rho)$ is quite cumbersome for a quantization scheme with many levels. However, here we are mainly interested in the loss of SNR for a signal much weaker than the average noise, i.e. when ρ is small. Linearizing the binormal distribution in ρ, one obtains that the efficiency of a given quantization scheme, defined by

$$E = \frac{1}{\rho} \frac{< X_1 X_2 >}{< X_1^2 X_2^2 >^{1/2}} \tag{2}$$

can be computed analytically.

Let $X(x)$ defined by a set of thresholds l_i and corresponding values v_i (all positive, $i = 1, \ldots, n$), in such a way that $X = v_i$ if $x \in [l_i, l_{i+1}]$. The first threshold is $l_0 = 0$, and threshold $l_{n+1} = \infty$. If one imposes that $X(x)$ is odd, a quantization scheme with an even number of levels ($2n$) is defined. For a scheme with an odd number of distinct levels, one can set $v_1 = 0$.

Usually the correlator computes the products $X_1 X_2$ by a multiplication table, Vij. It is not required that $V_{ij} = v_i v_j$, even if this is the simplest assumption.

By computing explicitly E, and exploiting the symmetry of V_{ij}, one obtains

$$E = \frac{2}{\rho} \frac{\sum_i \int_{l_i}^{l_{i+1}} V_{ij} P(x_i, x_j; \rho) dx_i dx_j}{\left(\sum_i \int_{l_i}^{l_{i+1}} V_{ij}^2 P(x_i, x_j; \rho) dx_i dx_j\right)^{1/2}} \tag{3}$$

Consistently with the above definition, here and in all subsequent formulas the summations are taken over positive values of i only.

Linearizing $P(x_1, x_2; \rho)$ one obtains $P(x_1, x_2; \rho) = P(x_1, x_2; 0)(1 + x_1 x_2 \rho)$, and

$$\begin{aligned} E &= 2 \frac{\sum_i V_{ij} Q_i Q_j}{\left(\sum_i V_{ij}^2 R_i R_j\right)^{1/2}} \\ Q_i &= \frac{1}{\sqrt{2\pi}} \int_{l_i}^{l_{i+1}} x \exp\left(-\frac{x^2}{2}\right) dx \end{aligned}$$

$$
\begin{aligned}
&= \frac{1}{\sqrt{2\pi}}\left(\exp\left(\frac{-l_{i+1}^2}{2}\right) - \exp\left(\frac{-l_i^2}{2}\right)\right)\\
R_i &= \frac{1}{\sqrt{2\pi}}\int_{l_i}^{l_{i+1}} \exp\left(-\frac{x^2}{2}\right) dx\\
&= \frac{1}{2}\left(\mathrm{erf}\left(\frac{l_{i+1}}{\sqrt{2}}\right) - \mathrm{erf}\left(\frac{l_i}{\sqrt{2}}\right)\right) \qquad (4)
\end{aligned}
$$

The above expression can be used to maximize E with given constrains, e.g. equally spaced thresholds, or integer v_i. In table 1 the quantization loss for several equally spaced, 8 level sampling schemes, is listed. Sampling levels are integer multiples of the quoted quantization step, always optimized for maximum efficiency. For comparison, the best quantization scheme with 8 levels has a quantization loss of 3.45%.

Threshold step	Weights	Quantization loss (%)
0.565	1, 3, 5, 7	3.74
0.555	1, 3, 5, 8	3.62
0.565	1, 3, 5, 7.66	3.58

Table 1: Quantization losses for various fixed step, 8 level sampling schemes. Thresholds are expressed for a normal Gaussian distribution. The first case is for uniformly spaced level values, the second for best integer values, and the third for minimum quantization losses. The best quantization scheme for nonuniformly spaced 8 level quantization has a loss of 3.45%

The expression for E can be derived in a more intuitive, albeit less rigorous, way. If (as in our case) $V_{ij} = v_i v_j$, summations in the expression (4) separate into two identical parts, and $E = <X>^2 \,/\, <X^2>$. In other words, the efficiency can be computed from the mean and variance of the two sampled signals, simpler to compute than the correlation product. If x is composed by pure uncorrelated Gaussian noise plus some small signal $\delta x(t)$, then $<X(t)>$ may be thought as the nonzero average due to the signal. $<X^2>$ is the variance of the digitized noise only.

The expected value $<X>$ can be derived observing that the small offset in x is equivalent to slightly move the thresholds l_i. Therefore the contribution of each level v_i is equal to δx times the probability of x at the thresholds l_i: $<X> = 2\sum_i v_i Q_i$. This result is identical to that found in equation 4, but the method can be applied to compute the quantization loss of more complex processes. We will use it to compute the loss due to the quantization after the digital mixer.

For simplicity, we will consider the output of the imaginary channel only. The complex case is analogous, but mathematically more cumbersome. The efficiency will be computed comparing the result for the quantized signal to what is obtained in an analog (or infinite precision digital) case.

If the mixer output Y is quantized using a set of thresholds l'_i and values v'_i, its expected value $<Y>$ and the corresponding variance are:

$$
\begin{aligned}
<Y(t)> &= 2\delta x \sum_i Q_i Y(v_i \sin(\omega_0 t))\\
<Y^2> &= 2\sum_i R_i Y(v_i \sin(\omega_0 t))^2 \qquad (5)
\end{aligned}
$$

The summations are always carried on positive levels only. $Y(x)$ is the quantized value for x (see fig. 6).

$<Y(t)>$ still depends on the phase of the LO. To compute the amplitude of the sinusoid, the Fourier component at the LO frequency $<Y(t)> sin(\omega_0 t)$ must be evaluated. These values must be averaged over a cycle of the digital LO phase. Assuming that the cycle is divided into N bins,

$$\begin{aligned} < Y > &= \frac{\sqrt{2}}{N}\delta x \sum_i Q_i \sum_j \sin(2\pi j/N) Y(v_i \sin(2\pi j/N)) \\ < Y^2 > &= \frac{2}{N}\sum_i R_i \sum_j Y(v_i \sin(2\pi j/N))^2 \end{aligned} \tag{6}$$

For an infinite precision digital mixer, we obtain that $< Y >^2 / < Y^2 >= 0.5$ Therefore the total loss in SNR due to *both* quantizations is

$$E = 2\frac{< Y >^2}{< Y^2 >} \tag{7}$$

The set $\{l', v'\}$ can be determined by maximizing $< Y^2 > / < Y^2 >$. Using the input quantization scheme listed in the last row of Table 1, and various quantization weights for the mixer output, we found the thresholds that minimize the total quantization loss. The additional loss due to the sampler is comparable to that of the first quantization. However, it is more important to minimize the spur levels due to the harmonics of $< Y(t) >$, as it will be shown in the next chapter.

6 Spur levels for a digital sinusoid

A digital BBC is more prone to unwanted conversion products with respect to a analog system, because there is no possibility to filter the LO or the mixer output. Unwanted conversion products are aliased in band, and high order terms may easily been generated.

The system can be schematized as in fig. 6. The input signal $x(t)$ is first converted to a few-bit digital representation $X(t)$. The oscillator output is sampled at an appropriate multiple of its natural frequency, and convolved with a rectangular window, corresponding to the phase bin of the tabulated representation. The resulting function is resampled at the correlator sampling frequency (4 GHz) and multiplied with the digital signal. After multiplication (in principle, with infinite precision), the result is re-quantized to the signal $Y(t)$, that is sent to the digital filters.

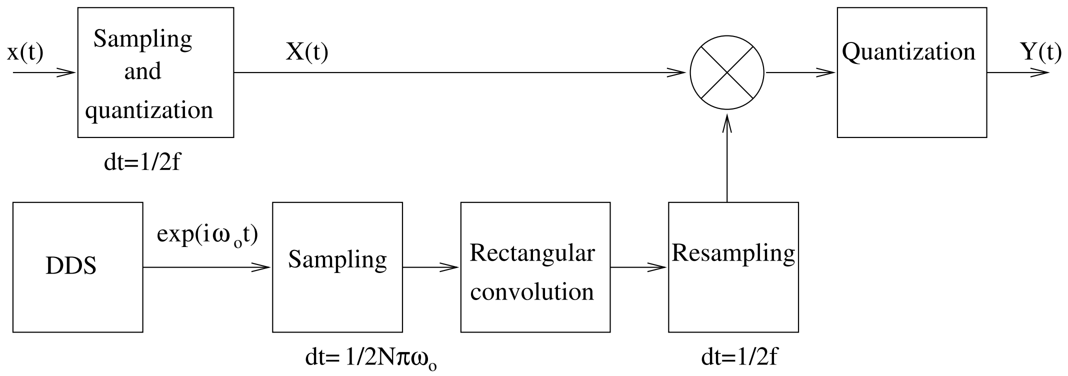

Figure 6: Model used to analyze the performance of a digital LO/mixer. The sinusoid generated by an ideal oscillator is sampled at the frequency $N\omega_0$, convolved with a rectangular window of width $\tau = 2\pi/N\omega_0$, multiplied with the sampled digital signal, and then resampled.

Two steps introduce the possibility of undesired spurs. Quantization and windowing of the sinusoid produce harmonics of the LO signal, with an order equal to the number of phase bins. Quantization of the result is a nonlinear operation, and produce harmonics of almost everything.

LO harmonics produce undesired conversions if strong lines are present together with weak ones. Since usually astronomical lines are not large compared to the noise level, the most dangerous spurious lines are of the form $nf_{lo}+f_{in}$. Amplitude of these products depend linearly on input line level, i.e. the ratio between the strongest input line and the spurious products is constant. Numerical simulations have shown that this is the case even when the strongest line has a monochromatic power of a few percent of the total input power.

LO harmonic level can be computed analytically. If the LO is represented with a given number N of frequency bins, each bin has a duration $\tau = 2\pi/N\omega_0$. Its spectrum has harmonics at frequencies $\omega_k = \omega_0(1+kN)$, with k integer. Amplitude of term k is given by

$$A_k = \mathrm{sinc}\left(\frac{\omega_k \tau}{2}\right) = \mathrm{sinc}\left(\frac{\pi}{N} + k\pi\right) \simeq \frac{1}{1+Nk} \tag{8}$$

Therefore $A_k^2 \simeq (Nk)^{-2}$. For $N = 64$ and $N = 128$, the first harmonic is resp. 36 dB and 42 dB below the carrier. These harmonics are aliased in the used bandwidth by the resampling at 4 GHz. It is not possible, as in an analog system, to filter them, because the sinusoid generation, windowing, resampling and mixing are done atomically in the RAM LUT's.

The effect of the second quantization is analyzed by computing the response of the overall system (sampler, first quantization, mixer and second quantization) to a small signal superimposed to a Gaussian noise. When a small DC input is applied to the sampler input, the mixer output has an expected value, syncronous to the LO phase, that can be computed using equations (5). The response of the system to small signals may be considered as the product of the input signal to a digitized approximate sinusoid computed in this way. This digitized sinusoid has only even components, since the quantization function Y is odd. If $Y(\omega)$ is the Fourier transform of $< Y(t) >$, we can reduce the harmonic content by minimizing the sum of $|Y(\omega/Y(\omega_0)|^k$, with k a positive integer. A high value for k favour a solution with relatively high harmonics, but with a minimum value for the higher one.

The minimization is performed on one channel only (the imaginary one). It can be shown that the complex LO has the same harmonic content of the real part of one of the two channels, but with harmonics of the form $4n-1$ converted to negative frequencies.

The minimization problem is ill conditioned, since the quantization function $Y(x)$ is nonlinear and is evaluated over a finite number N of points. Therefore, the dependance of the quality function on many parameters (e.g. the quantization thresholds l'_i) is stepwise, and most minimization algorithms would fail.

An algorithm that does not use derivates is the simplex method, described by Nelder and Mead [3], and implemented e.g. in *Numerical Recipes* [2]. The harmonic content, computed with $k = 4$, has been minimized by varying the thresholds of the second quantization, l'_i. The first quantization has been kept fixed with constant threshold spacing and weight values $v_i = \{1, 3, 5, 7.66\}$. The values $\{1,\ 3,\ 5,\ 8\}$, more appropriate for the subsequent digital processing, have been chosen for the second quantization, v'_i.

The algorithm found a solution with the thresholds listed in Table 2. An optimal solution for spur levels has a quantization loss around 9% (including a 3.65% due to the first quantization of $X(t)$), and the highest harmonics have an energy content of -31.3 dBc. A suboptimal solution for spur levels have a quantization loss of 6.9%, slightly above the minimum for a 8 level scheme. A spurious conversion of -30 dB is problematic for spectral line observations. It is not uncommon to find strong lines, well above 1000 times the noise floor. These lines will produce significant ghosts, at unpredictable frequencies, that can be discriminated from real lines only by displacing the LO by a small value.

A much better approximation can be obtained if the second quantization has 4 bits per sample. This has however a significant impact for the digital filter design, since input/output pins and board space for traces are increased by 33%. The spur level is reduced to -38.4 dBc, that becomes probably acceptable for most spectroscopic requirements. The quantization

Mixer Thresholds	Mixer Weights	Quantization loss (%)	Max. spur level (dBc)
0.436, 0.869, 1.279	1, 3, 5, 7	6.90	-23.0
0.446, 0.885, 1.446	1, 3, 5, 8	6.54	-27.3
0.453, 0.867, 1.618	1, 3, 5, 8	6.90	-31.2
0.504, 0.906, 1.094	1, 3, 5, 8	9.07	-31.3

Table 2: Quantization losses and spur level for a optimal weight, uniform threshold sampler followed by a digital mixer. Quantization loss refers to both input and mixer quantizations. Various quantization schemes for the mixed signal are examined. The first two cases are for minimum loss, the second two for minimum spur levels.

loss due to the mixer drops to 1.5% (total loss of 5.10%). Spur level depends somehow on sampler threshold accuracy. The dependance is not very strong, a variation of even 20% on each threshold degrade spur level of at most 1 dB.

In fig. 8 the output of the digital mixer is plotted, both for a 3-bit and a 4-bit mixer. Solid lines represent the mixer output when the mixer input has one of the possible quantized values produced by the sampler. These waveforms are those actually stored in the RAM tables. The dotted line represents the mixer output for a small signal embedded in Gaussian noise, as given by equation (5). The mixer harmonic content has been evaluated on this waveform.

It should be noted that the considerations expressed in this chapter affect any design in which a digital LO is employed, independently on whether and how a digital correlation scheme is used.

7 Simulations

A set of programs have been written to simulate all elements of the proposed correlator. Both continuous ad digitized versions of each element are available, to evaluate independently quantization errors.

A block diagram of the simulation software bench is shown in fig. 7.

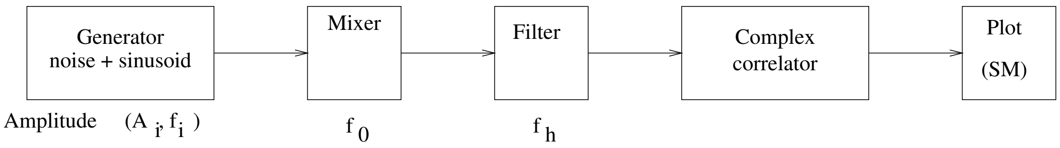

Figure 7: Simulation chain for the digital BBC

Each element is represented by a program that acts as a Unix filter. The whole simulation is executed as a pipeline of Unix programs. Single elements can be substituted with different versions (e.g. digital or analog), or omitted.

A generator produces pseudorandom noise, with an arbitrary number of lines added. The signal is then processed by a mixer (that includes digital sampling/quantization), optionally filtered, and analyzed by a complex correlation spectrometer. Since the generator is deterministic, it is possible to reduce the noise in the output spectrum by subtracting a noise only reference spectrum. The simulation engine processes from 2000 to 10000 samples/s on a Sparc workstation.

In figures 9 and 10 an artificial spectrum is shown during various processing stages. The original spectrum is shown in fig. 9a. The negative portion of the spectrum is also sohown, even if it is just the mirror image of the positive portion and then carries no further information, because it will become important in the successive stages of signal processing.

This spectrum is composed of Gaussian noise, with a bandwidth of $0.5f_c$, plus three monochromatic lines. All frequencies are expressed as a fraction of the sampling frequency, f_c. Total simulation involves 10^7 samples, giving an espected fractional noise of about 1%. The line frequencies are 0.1, 0.2 and $0.31f_c$, and their amplitudes are respectively 0.07, 0.07 and 0.1 times the noise RMS amplitude σ. The LO frequency and the final bandwidth are plotted for reference. The line peak amplitude in the plotted spectra may vary a little due to frequency resolution effects, and plotted line ratios differ somewhat from the values listed above.

In fig. 9b we can see the spectrum of the signal after conversion, with a LO frequency of $0.27f_c$, and re-quantization. The frequency scale has been rotated, and now frequency 0 corresponds to the frequency of the digital LO ($0.27f_c$, in this case). The thresholds adopted for the sampler are those for optimum 3-bit, uniform spaced levels, reported in the last line of table 1. The thresholds for the second quantization, (and consequently the tabular representation of the digital sinusoid) are those listed in the third row of table 2. Now positive and negative frequencies carry different informations.

The signal is then filtered, using a digital finite impulse response (FIR) low-pass filter, with a cutoff frequency of $f_c/8$ and a total length of 192 taps. The filter shape has been chosen as a truncated and tapered $\sin(x)/x$, without a particular care for accurate response (see [9] for computation of accurate filters). The resulting spectrum is shown in fig. 9c. The filter output is then decimated (resampled) by a factor of 4, expanding the frequency scale by the same factor (fig. 10). The total bandwidth is $0.25f_c$, i.e. half the bandwidth of the original signal. The effective bandwidth is divided among positive and negative frequencies.

To analyze the performance of the BBC with respect to the spurious conversions, a spectrum with strong lines has been processed. Line intensity has been increased to 0.1, 0.1 and 0.2 σ, respectively, with the same frequencies of the previous case. A total of 10^8 samples have been processed. A reference spectrum, with only the same pseudorandom noise sequence and no lines, and processed in the same way, has been subtracted, to further reduce the noise floor by about a factor of 4. The result, expressed in a logarithmic scale, is shown in fig. 11. This plot is basically a logarithmic scale version of fig. 9b. First spurious signals can be seen at a level of approximately -32 dBc. Noise level is around -35 dBc, and it is not easy to discriminate between spurious conversions and noise peaks, but an upper limit in reasonable agreement with the considerations of chapter 6 is verified.

8 Conclusions

A digital BBC appears feasible from a technical point of view, and does not increase significantly the complexity of the proposed ALMA correlator. Due to the digital LO can be implemented with a pair of FPGA's placed between the sampler and the digital filter.

Several approaches for sideband rejection are possible, the simplest ones being a differential phasing network included in the digital filters, and a complex correlator implemented using the memory buffer board.

The extra quantization required by the digital LO introduces a SNR loss of a few percent. The digital nature of the LO, however, high order armonics of the LO frequency are produced, and aliased in band. Strong lines beat with the LO armonics, creating ghost images. A rejection of unwanted conversion products of approximately 30 dB can be achieved with a 8-level quantization of the mixer output. To increase rejection to 40 dB, a 16-level quantization is required, increasing the data path from the mixer to the digital filters to 4 bit (128 signals for a 32 times multiplexed signal).

Italian participation to the ALMA project is supported by Italian CNAA.

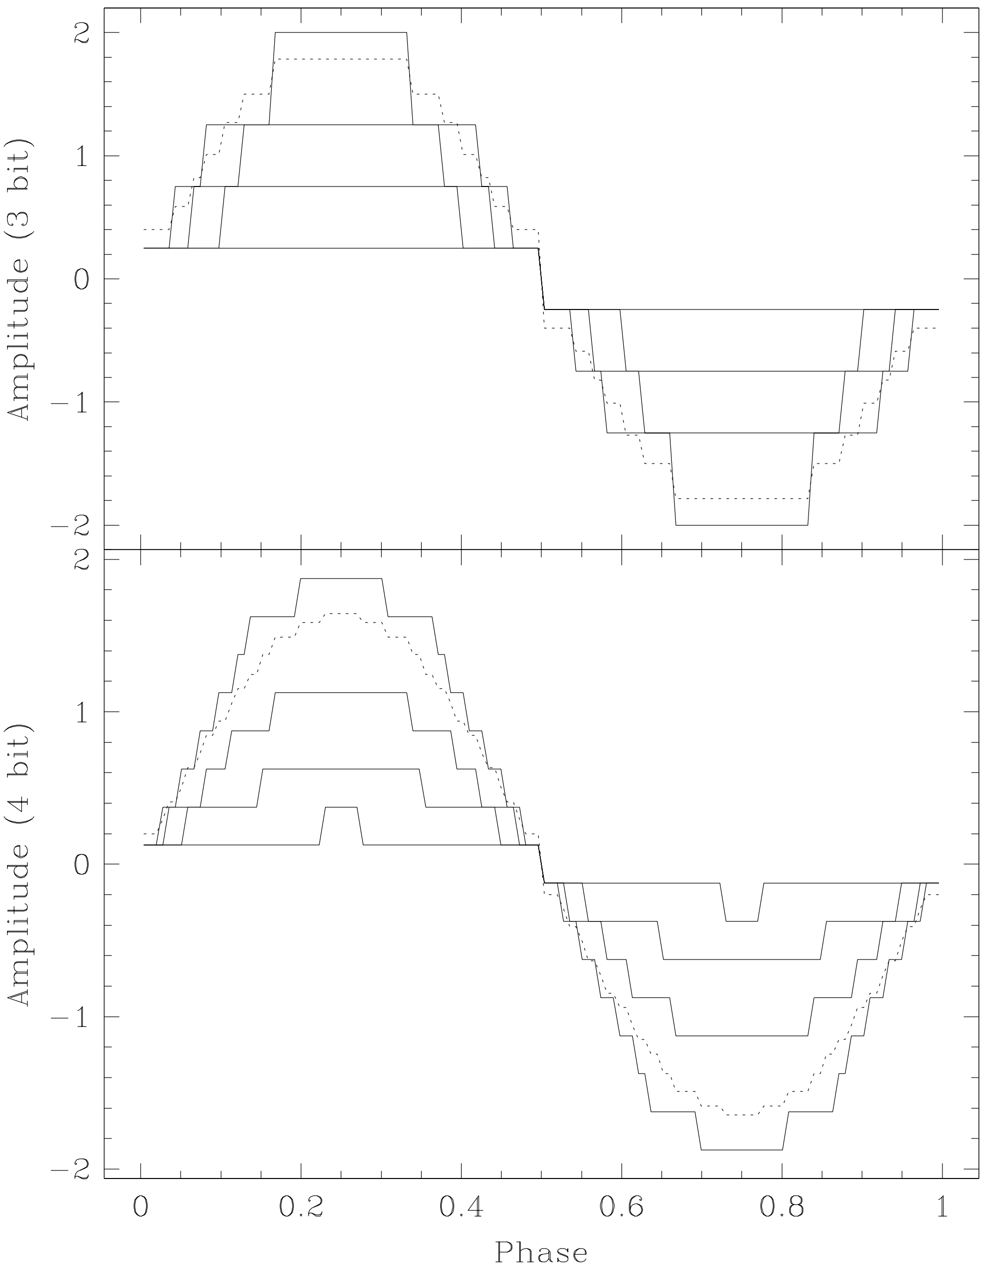

Figure 8: Synthesized sinusoid for the digital LO, resp. for 3 and 4 bit mixer, and a 3 bit sampler. Solid lines represent quantized sinusoid for each value of the sampled signal. Dotted line represents weighted output for a Gaussian sampler input

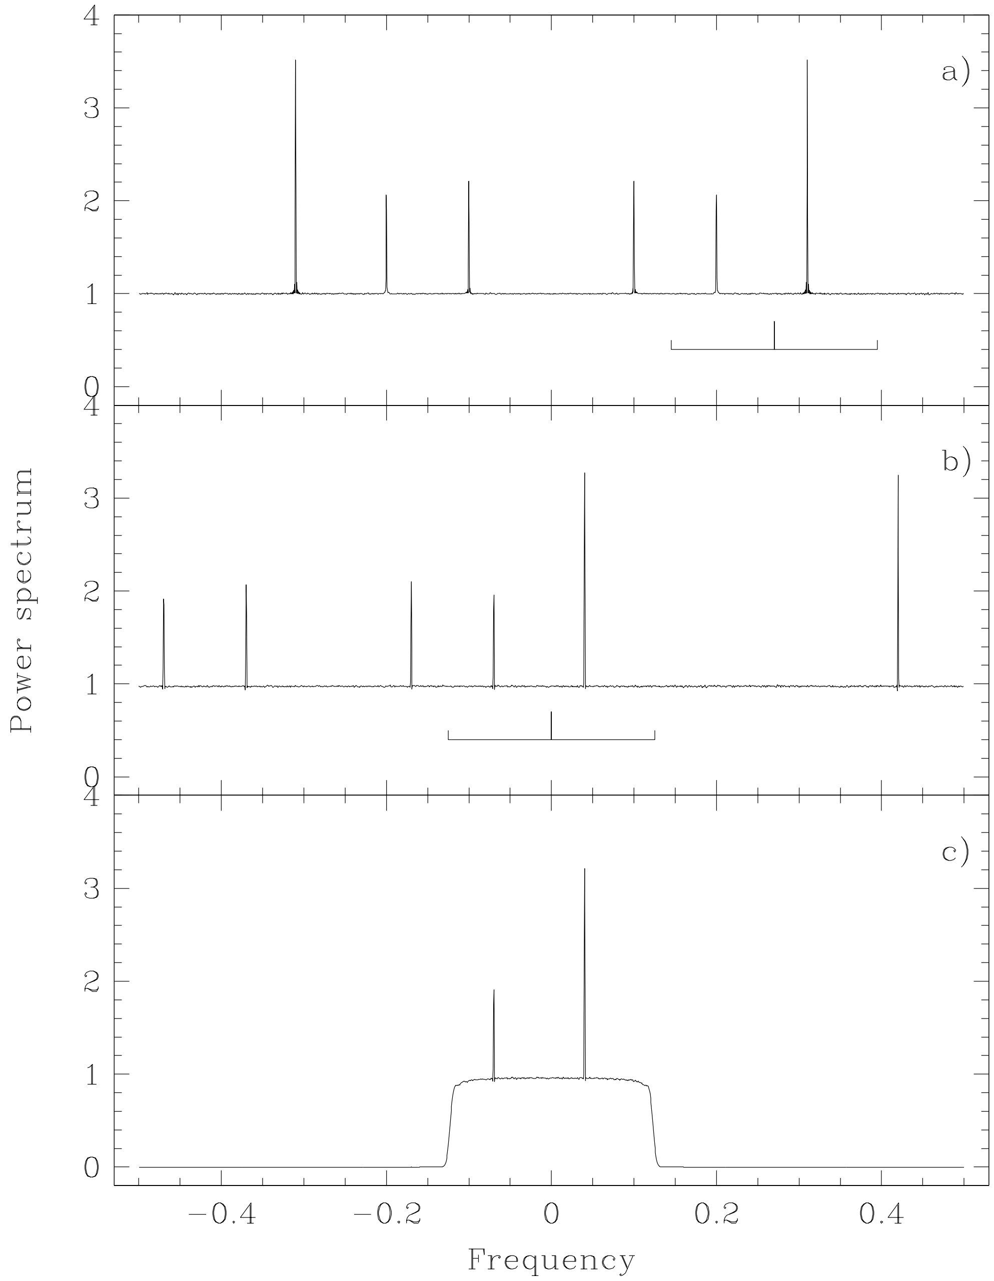

Figure 9: Processing of a test spectrum for the BBC simulation. It is composed of Gaussian noise, and 3 lines (see text). The frequency scale is expressed in term of the sampling frequency. ¿From top: a) Original signal; b) after downconversion, with a LO frequency of $0.27 f_c$; c) After digital filtering, with a low pass FIR filter

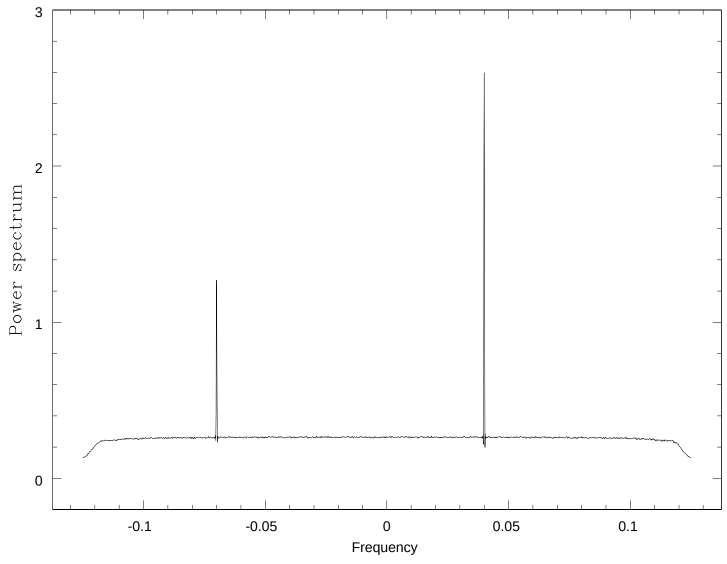

Figure 10: Spectrum after downconversion by the digital mixer, digital filtering and resampling at $f_c/4$.

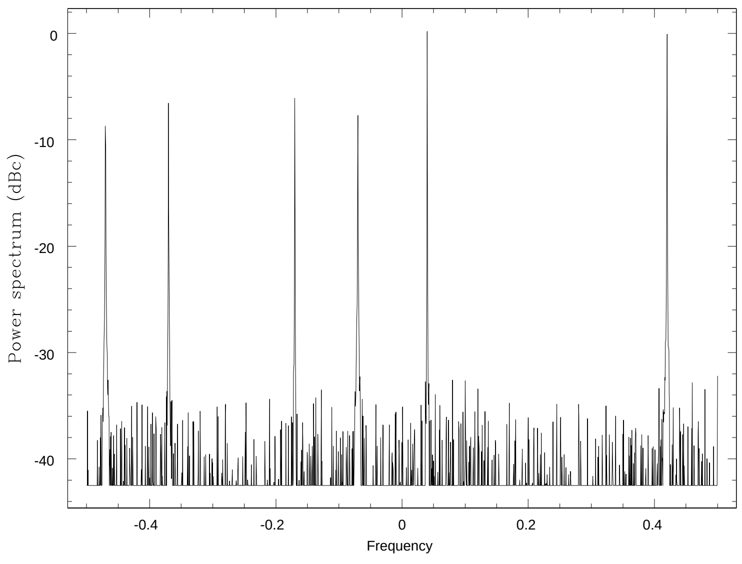

Figure 11: Spectrum after downconversion by the digital mixer, on a logarithmic scale. Spurs are visible at -32 dBc

References

[1] Cooper, B.F.C.: *Correlation with Two-bit Quantization*, Ausl. J. Phys, **23**, 521 (1970).

[2] Press, W.H., Flannery, B.P., Teukolsky, S.A., Vetterling, W.T.: *Numerical Recipes in C*, 306, Cambridge Univ. Press (1988)

[3] Nelder, J.A, Mead, R. Computer Journal, **7**, 308 (1965).

[4] Escoffier, R.: *The MMA Correlator*. MMA report **166** (1997).

[5] Thompson, A.R.: *A system design for the MMA*. MMA Report **190** (1997).

[6] Escoffier, R., Webber, J.: *Digital Filtering in the MMA*. MMA report **204** (1998).

[7] Thompson, A.R.: *Quantization Efficiency for Eight or more Sampling Levels*. MMA report **220** (1998).

[8] Escoffier, R., D'Addardio, L.R., Broadwell, C., Webber, J., Schwab, F.: *Computer Simulation of a FIR Filter for the MMA*. MMA report **248** (1999).

[9] Anderson, B.: *FIR Filter Designs for ALMA*, web page `http://www.jb.man.ac.uk/~ba/firf.html` (1999).

Contents