

The new 3-stage, low dissipation digital filter of the ALMA Correlator

P.Camino¹, B. Quertier¹, A.Baudry¹, G.Comoretto², D.Dallet³

¹ *Observatoire de Bordeaux, LAB, Université de Bordeaux
BP 89, 2 rue de l'Observatoire, F- 33271 Floirac*

Email: camino@obs.u-bordeaux1.fr, quertier@obs.u-bordeaux1.fr, baudry@obs.u-bordeaux1.fr

² *Osservatorio Astrofisico di Arcetri, Largo Fermi 5, I- 50125 Firenze*

³ *Laboratoire IMS, ENSEIRB, Université de Bordeaux, 351 Cours de la Libération, F- 33405 Talence*

Abstract- The main goal of this study is to reduce the power dissipation of the 2-stage digital filter used in the ALMA Correlator system. This has been achieved by optimizing the number of FPGA logic elements used for the filter implementation. We have investigated the implementation of various structures based on the Cascaded Integrator Comb (CIC) filter in order to replace the present first filter stage, a 32-time demultiplexed input decimation filter. We conclude that a CIC filter cascaded with a quarter-band filter significantly improves the overall power dissipation and thus the FPGA thermal behaviour and reliability. This new design results in a significant improvement (nearly 25%) in the dissipation of each one of the ALMA filter cards.

1. INTRODUCTION

All the independent antenna pairs of the ALMA array are processed up to a maximum of 2016 by the ALMA baseline correlator system whose main specifications are given in [1]. In this memorandum, we recall the main functions of the digital filtering sub-system adopted for the ALMA baseline correlator and concentrate on the need to optimize the filter cards power dissipation (Section 2), we describe the solution proposed (Section 3), and present, after implementation of our design in the production filter cards, the results of our power consumption measurements (Section 4).

2. THE TUNABLE FILTER BANK (TFB) AND OPTIMIZATION OF POWER DISSIPATION

2.1. Electronic architecture

The main functions of the digital filtering sub-system of the baseline correlator were initially described in [2]. This sub-system is named Tunable Filter Bank (TFB) and is schematically shown in Fig. 1. The TFB specifications are given in [3]. In addition to the Direct Digital Synthesizer (DDS) which provides the 'tunable feature' of the TFB the filtering sub-system consists of two cascaded low-pass Finite Impulse Response (FIR) filter stages. FIR filters have a linear phase variation across the bandwidth as required for radio interferometry applications. The aim of this sub-system is to extract, by frequency division of the input signal bandwidth, sub-bands of smaller bandwidths in order to perform higher spectral resolution analysis. Multi-resolution analysis of different spectral regions is possible thus allowing optimal zooming of the most interesting spectral features. Analysing the full bandwidth in separate sub-bands results in increased spectral resolution.

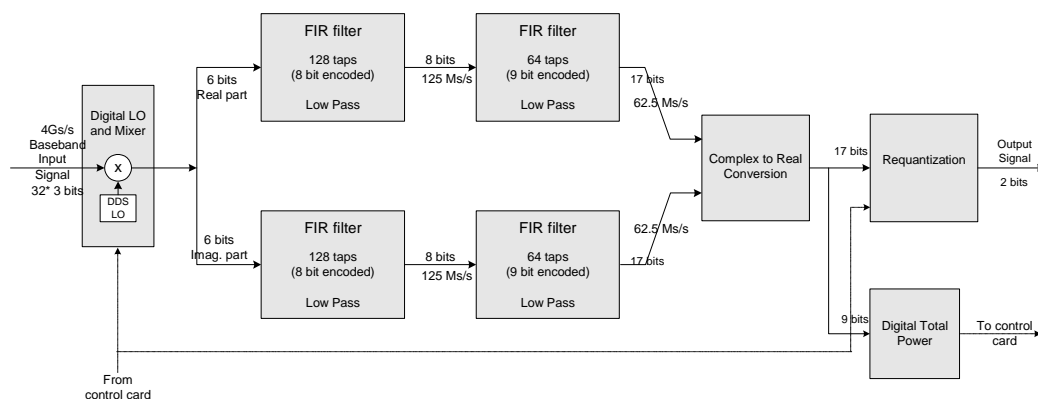


Fig. 1: Original TFB FIR architecture

The incoming signal is the wide ALMA baseband signal (2-4 GHz), digitized at 4 GS/s by a 3-bit, 8-level Analog to Digital Converter (ADC) specifically designed for the ALMA project [4]. The 4 GS/s input rate delivered by the ADCs cannot be processed by the FPGAs in the TFB as they are limited by their maximum clock frequency. To comply with the 125 MHz clock rate selected for the ALMA filtering and correlator sub-systems, the sampled signal is represented as 32 time-demultiplexed data lines at 125 MS/s with each line corresponding to one of the 32 successive samples of the digitised 2-4 GHz input data flow.

The frequency conversion required to select each sub-band position yields a complex signal. Real and imaginary parts of this signal (Fig. 1) are processed identically in time domain and later recombined to provide real samples to the correlator cards. From now only one data stream will be considered (Fig. 2). The first decimation filter has broad transition band specifications, and a passband of $1/32^{\text{th}}$ of the input band. Attenuation in the stop-band is -47 dB and the passband ripple is 0.2 dB. It is followed by a decimation in time process with a decimation factor of 32. The second decimation filter stage is a half-band filter with a decimation factor of 2 (Fig. 2). The final attenuation and passband are combination of the 2-stage filter cascade but the final transition region is fixed only by the second stage.




Fig. 2: Data processing

To cover the entire 2 GHz input band and to meet the sampling Nyquist conditions, 32 sub-band filters are implemented, each one synthesizing a bandwidth of 62.5 MHz. A total of 512 TFB cards are required for the complete Correlator System. Each card is populated with 16 FPGA chips (2 sub-band filters per FPGA, 90 nm technology).

2.2. Power consumption and junction temperature issues

Based on the architecture shown in Fig. 1 our original FIR filter design has been implemented in Stratix II chips from Altera. We measured a total dissipation of about 78 W per card. Despite this improvement from the 100W ALMA Correlator specification, the chip junction temperature expected at 5000-m elevation in the operational conditions of the complete Correlator System would remain close to the maximum temperature recommended by Altera. This would negatively impact the expected failure rate, resulting in significant maintenance problems. Therefore, Correlator IPT decided to lead two actions in parallel: a) to improve the air flow circulation in the Correlator Station racks, b) to consider how a redesign of the TFB first-stage filter could improve the dissipation per filter card, and implement the alternative design.

3. A NEW DESIGN BASED ON THE CIC FILTER PRINCIPLE

In this Section we describe the new FPGA personality developed for the TFB sub-system with the main goal of reducing the TFB card dissipation. The original 2-stage filter structure has been replaced by a 3-stage filter structure (see Fig. 3) based on the use of a Cascaded Integrator Comb filter (CIC).

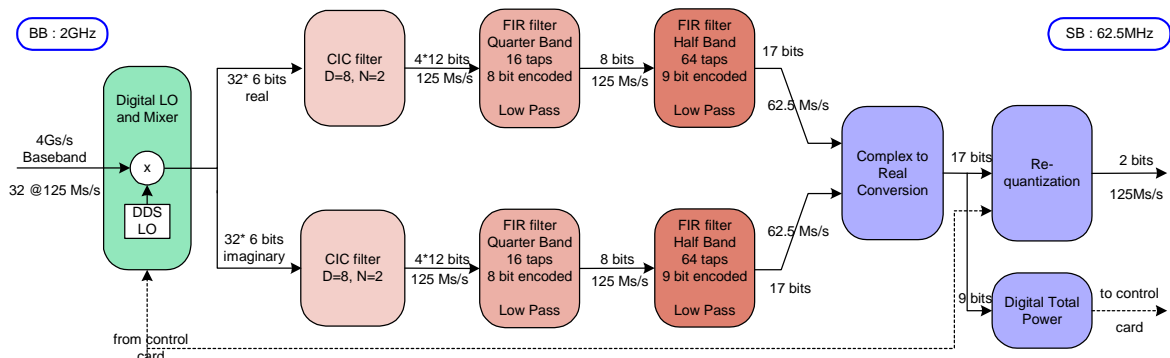


Fig. 3: Electronic architecture of the 3-stage TFB

3.1. A multi-stage filter

The main idea for decreasing the TFB power consumption consists in finding alternative designs to the original first filter stage which uses the largest share of logic elements (Adaptive Logic Modules in ALTERA Stratix II designation). The distribution of resources for the original design is given in Table 1.

Table 1. Distribution of resources for one TFB filter

	ALMs	Mem. 512 bits	Mem. 4k bits	M-RAM	Mult. 9*9 bits
DDS	1191	1	57	/	/
1 st filter stage	1680*2	/	/	/	/
2 nd filter stage	213*2	31*2	/	/	16*2
Conv/Requant stage	84	1	/	/	2
Misc.	1123	4	14	/	2
Overall (2 TFBs)	10970 (81%)	96 (48%)	103 (72%)	0 (0%)	70 (55%)

Being the first filter a decimation filter with a large transition region, we considered a CIC filter solution. The CIC transfer function is given by [5]:

$$H(z) = \left(\sum_{k=0}^{D-1} z^{-k} \right)^N = \left(\frac{1-z^{-D}}{1-z^{-1}} \right)^N \quad (1)$$

The most interesting feature of this kind of filter is the unitary format of its taps. The second representation of (1) is a sum transformation resulting in an Integrator part and Comb part cascade (classical implementation). The CIC transfer function is fully defined by the decimation factor D and the filter order N (Fig. 4(a)).

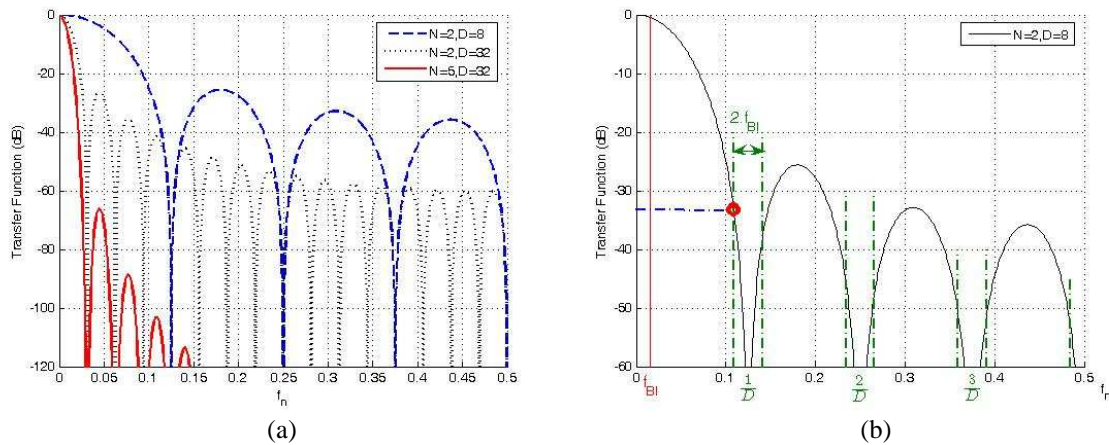


Fig. 4: CIC Magnitude frequency response

High out-of-band selectivity can be obtained with relatively low order compared to other filter structures. Note that increasing the order results in a faster passband drop which will have to be compensated in a second filter. The linear phase characteristic across each magnitude response lobe has to be pointed out (the CIC filter is a Finite Impulse Response filter). For a defined final passband f_{BI} (in the normalized frequency space) and different decimation factors D , we can use the attenuation table shown in Table 2 to determine the optimal value for N and D . The decimation is expressed as $f_{BID}=D \cdot f_{BI}$, and the attenuation is computed at frequency $f_c=1/D - f_{BI}$, i.e. the frequency where the worst case of aliasing error occurs, as illustrated by red point on Fig. 4(b) for $f_{BI}=1/64$; note that the green dotted lines delineate the regions folded in the passband after decimation.

Table 2. Worst case attenuation (at f_c) for different values of D and N

	$N=1$	$N=2$	$N=3$	$N=4$	$N=5$
$f_{BID}=1/4$	-10 dB	-20 dB	-31 dB	-42 dB	-52 dB
$f_{BID}=1/8$	-17 dB	-34 dB	-51 dB	-68 dB	-84 dB
$f_{BID}=1/16$	-23 dB	-47 dB	-70 dB	-93 dB	-116 dB
$f_{BID}=1/32$	-28 dB	-58 dB	-86 dB	-115 dB	-144 dB
$f_{BID}=1/64$	-35 dB	-71 dB	-105 dB	-140 dB	-175 dB

Several implementation solutions have been investigated and are described in [6]. Due to the demultiplexed input format, the most appropriate CIC filter parameters, i.e. the highest decimation factor coupled with the lowest CIC filter order satisfying the attenuation specifications, are for the ALMA case ($f_{BI} = 1/128$): $D=8$ (namely $f_{BI} = 1/16$) and $N=2$. To achieve the required 1st TFB stage decimation factor of 32, we thus have to cascade to the CIC filter another filter allowing a decimation factor of 4. A quarter band FIR filter has been chosen (see Fig. 3).

3.2. Multi-stage electronic architecture

3.2.1. The CIC filter

Because no multipliers are required and no coefficient storage is needed for a CIC filter we expect a relatively easy implementation and low power dissipation in the filtering sub-system. The main electronic structures that can be found in the literature for CIC filters have been examined [6]. They include a classical structure, a modified rotated-angle CIC filter structure [7], a CIC polyphase decomposition, a non-recursive demultiplexed CIC filter structure, and a non-recursive CIC filter structure [8].

Due to the specific 32-time demultiplexed input format the optimal CIC implementation is a parallel non-recursive architecture. The transfer function of such a structure, after a factorization of equation (1) is given by:

$$H(z) = \prod_{i=0}^{(\log_2 D)-1} (1 + z^{-2i})^N \quad (2)$$

The schematic of the $D=8, N=2$ CIC filter is given in Fig. 5.




Fig. 5: Non recursive architecture ($D=8, N=2$)

Arithmetic operations are performed with full scale representation. Each block is followed by a decimation by 2 that allows us to suppress every other addition at the output of the blocks. Note that no signal truncation is performed at the CIC filter output; the output is encoded on 12 bits.

3.2.2. The quarter-band FIR filter

The FIR filter used to achieve the final decimation is a quarter-band (QB) filter with a large transition band. The transfer function of the QB filter is shown in Fig. 6: $[f_1, f_2]$ is the transition band and $[0, f_1]$ is the final band selected by the final TFB filter stage (note the passband drop). This filter has been synthesized with the Remez algorithm. It results in a 16-order quarter band FIR filter with symmetrical impulse response.

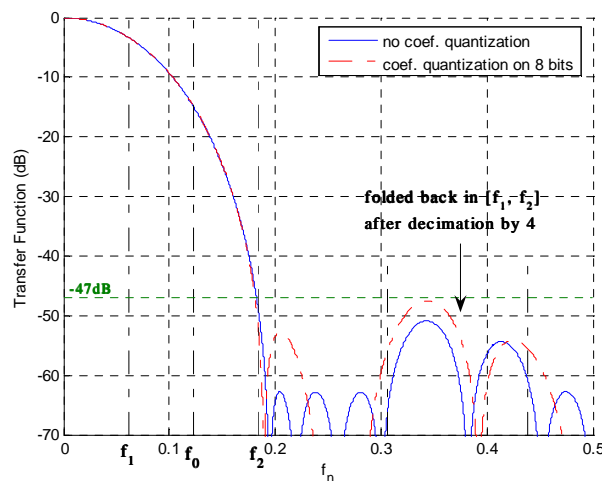


Fig. 6: Quarter-band filter transfer function. f_0 is the Nyquist frequency after decimation

The structure of the quarter-band filter is depicted in Fig.7.




Fig. 7: Implementation of QB filter

The shift register outputs corresponding to symmetric taps are summed together before each multiplier to optimize for the symmetric coefficients. The decimation process by 4 is intrinsic to the architecture. The filter output is truncated to 8 bits to fit the final TFB filter stage input range (see Fig. 3).

3.3. Results

This multi-stage filter (CIC-QB filter cascade) results in an optimization of the resources available in the FPGAs. Table 3 gives an overview of the required ALMs and maximum frequency achieved.

Table 3. Summary of the different studied solutions

	Resources (Stratix II)	Max. Frequency
TFB 1 st stage (original design)	1775 ALMs	180 MHz
non-recur. CIC (D=8,N=2)+QB	630 ALMs	200 MHz

The number of required ALMs is decreased by almost a factor of 3 compared to the original design and the 125MHz correlator clock rate is easily met.

3.4. The final filter stage

As in the original design the half-band FIR filter determines the final band characteristics. It also compensates the passband drop of the CIC-QB filter cascade. To synthesize such a filter, the output from the Remez algorithm has been fitted to the requested passband response using a minimization algorithm (simplex minimizer [9]). Fig. 8 shows the transfer function of the final stage.




Fig. 8: Final stage transfer function

The electronic architecture used to implement this final stage is described in [2].

4. VALIDATION AND POWER DISSIPATION OF THE NEW FILTER DESIGN

The multi-stage designed - CIC filter, quarter-band FIR filter, half-band FIR filter - has been first validated with a VHDL simulation tool (Modelsim) using input and output test vectors generated from a mathematical model.

Then the firmware has been implemented in the TFB card chips to perform full functional validation with the ALMA Test Fixture. Fig. 9 shows 2 adjacent sub-band spectra obtained with the Test Fixture, each one being 62.5 MHz wide.




Fig. 9: 2 adjacent sub-bands showing the auto-correlation spectra in dB across 62.5 MHz bandwidth (one sub-band of the 2 GHz input band)

One spectral line has been placed in each sub-band, no folded line appears. A flat passband is also obtained.

Power consumption measurements of TFB cards for both the original and the new filter designs have been made in the laboratory with a 200 feet/min air flow. The original filter design gives an average dissipation of 78W per card while the new design yields a total of slightly less than 60W, giving a power consumption improvement of nearly 25%. These dissipations correspond to the case where all 32 sub-bands are being used. The other important point, linked to the FIR chip lifetime, is the chip junction temperature which has been measured for different air flow values as shown in Fig. 10. The blue and pink curves correspond to the new and original filter designs, respectively.




Fig.10: TFB chip temperature versus air cooling

Both plots given here are an average of the junction temperatures measured in the 16 Stratix II chips assembled on a TFB card. There is actually a slight temperature gradient across the matrix of 4 by 4 FPGAs. With the system of fans installed in the Station racks where the TFB cards are being operated on site, the air flow should reach about 200 feet/min which corresponds to a junction temperature of around 62°C (almost 12°C improvement compared to the previous design). The FPGA junction temperatures for all cards in all bins of each rack will remain below the maximum temperature recommended by Altera thus enhancing the FPGAs average time between failures, the system reliability and significantly improving the overall power consumption of the Correlator system.

5. CONCLUSION

Several architectures based on the CIC filter have been considered to optimize the ALMA TFB filter power consumption. The main problem encountered during this work was the parallel input data flow, not suited for the classical CIC filter structure. We have shown that a non-recursive CIC structure followed by a quarter-band filter optimizes the overall dissipation by making optimum use of the FPGA ALM resources. After a validation phase demonstrating that all filter cards specifications are met, we have implemented this new design in the ALMA TFB cards. The net result is a nearly 25% improvement in the power dissipation of each TFB card compared to our original design thus providing enhanced lifetime of the FIR chips and improved use of the power for the overall correlator system.

References

- [1] R. Escoffier, J. Webber and A. Baudry, "64 Antenna Correlator Specifications and Requirements", ALMA System Document, 2005.
<http://edm.alma.cl/forums/alma/dispatch.cgi/documents/showFile/100591/d20050708085722/No/ALMA-60.00.00.00-001-B-SPE.pdf>
- [2] B. Quertier, G. Comoretto, A. Baudry, A. Gunst, A. Bos, "Enhancing the Baseline ALMA Correlator Performances with the Second Generation Correlator Digital Filter System", ALMA Memo, n°476, 2003
- [3] A. Baudry, P. Cais, G. Comoretto, B. Quertier, "Production Tunable Filter Bank Card, Technical Specification", Internal ALMA document, CORL-60.01.07.05-004-B-SPE
- [4] C. Recoquillon, A. Baudry, J.B Begueret, S. Gauffre, G. Montignac, "The ALMA 3-bit 4 Gsample/s, 2-4 GHz Input Bandwidth, Flash Analog-to-Digital Converter", ALMA Memo, n°532, 2005.
- [5] Eugene B. Hogenauer, "An Economical Class of Digital Filters for Decimation and Interpolation", IEEE Transaction on Acoustics, Speech, and Signal Processing, vol. ASSP-29, pp. 155-162, 1981.
- [6] P. Camino, "Etude Comparative de Diverses Structures de Filtres Numériques. Application aux Signaux à tres Large Bande et au Corrélateur ALMA", PhD Thesis, Université de Bordeaux, 2008.
- [7] F. Daneshgaran, M. Laddomada, "A Novel Class of Decimation Filters for $\Sigma\Delta$ A/D Converters", Wireless Communications and Mobile Computing, vol. 2, pp. 867-882, 2002.
- [8] Y. Gao, L. Jia et al, "A Comparison Design of Comb Decimators for Sigma-Delta ADCs", Analog Integrated Circuits and Signal Processing, n°22, pp. 51-60, 1999.
- [9] W.H. Press, S.A. Teukolsky, W.T. Vetterling, B.P. Flannery, "Numerical Recipes: The Art of Scientific Computing, Third Edition", chapter 10.10, Cambridge University Press, ISBN: 978-0521880688, 2007