

Design of a FIR filter
using a Virtex2 Xilinx chip

G. Comoretto¹

¹Osservatorio Astrofisico di Arcetri

Arcetri Technical Report N° 5/2002
Firenze, September 2002

Abstract

In the hybrid correlator proposed for ALMA, a large fraction of the total complexity and cost is represented by the digital filterbank. In this report, an alternative design for the filter unit is presented. This includes a digital baseband converter, to select an arbitrary portion of the input band, and a two-stage filter. The baseband converter allows the selection of a 62.5 MHz portion of the input band arbitrarily positioned in the 2 GHz IF bandwidth. In this way, it is possible to cover it with partially overlapped subbands, allowing a better control of edge effects. The filter section is composed by a 64-tap coarse FIR filter, used to allow a 1:8 decimation of the input data, and a 128 tap sharp filter, equivalent to a 1024 tap filter operating at the original data rate. The whole filter can be fit in a Xilinx XC2V1000 FPGA, with a reduction of a factor 4 in total gate usage.

1 Introduction

In the hybrid correlator proposed for ALMA, a large fraction of the total logic and correlator cost is represented by the digital filter bank. In this architecture [1], denoted in this report as *basic FIR design*, 32 filters, with a number of taps of the order of 1024 each, are required to split the input bandwidth into the same number of parallel sub-bands. In each filter, the desired band is chosen using an appropriate filter shape, and then aliased to baseband by resampling the filtered data at a reduced clock rate. For flexibility, it should be possible to select any sub-band in any filter. Reduced bandwidth can be achieved by cascading filters together, at the expense of a reduced number of sub-bands.

No provision has been made for tuning the sub-bands, i.e. each sub-band must fall inside one of 32 equal slices of the input bandwidth. At reduced bandwidth, a similar restriction applies, i.e. the input bandwidth is always divided into N fixed sub-bands (with N a power of 2), and one individual sub-band is chosen among them. No provision is made for fine control of the individual sub-band phase, to implement a "fractional bit shift". Fractional bit shift is provided by adjusting the phase of the sampling clock, using a PLL to generate a small frequency offset in this clock.

In this report a slightly different approach is suggested, with a full digital SSB converter (in VLBA terminology, *Baseband Converter*, or BBC) to convert each sub-band to the required correlator input bandwidth. The circuit is composed of a digital local oscillator (LO), with enough resolution to provide "fractional bit shift" phase correction, followed by a two-stage FIR filter with (almost) fixed coefficients. This approach has the following advantages:

- fractional bit shift in sub-bands is easier to control than in the sampler, where resynchronization problems of the digital stream may arise with a variable-phase clock¹.
- Sub-band stitching can be done by using two extra "correlator slices" and packing the sub-bands closer to each other. This relaxes the requirements on the filter sharpness (wider transition band), and improves the phase response near the sub-band edge.
- Sub-bands can be positioned without any restriction in the input band, giving more observing flexibility. For example, it is possible to perform high resolution observations on several lines arbitrarily placed in the IF band.
- The possibility of placing the sub-band arbitrarily within the main IF band allows for more efficient filter design. In particular, the 2-stage FIR filter, described in this report, becomes possible, with considerable saving in total filter size and cost.

The digital BBC has some disadvantages with respect to the basic FIR filter design. These consist mainly in a worse SNR, due to the extra quantization in the mixer, and in possible intermodulation effects, because of the nonlinearity in the quantization process. However these effects can be controlled[6], and the overall signal degradation can be negligible (1.5% degradation in SNR). With respect to a WIDAR-type correlator[2], a digital LO placed before the bandpass filter has the disadvantage of producing intermodulations and ghost images in presence of strong input lines. These effects can be reduced to an acceptable level by using a more sophisticated, 3-bit representation for the LO sine wave, compared to the 2-bit representation that is sufficient in the WIDAR system.

2 Architecture

The system proposed here is composed of a bank of identical BBCs. Each BBC is composed of a digital LO, and a fixed filter. Filter shape is the same for all sub-bands, and need not to be changed to adjust bandwidth. The other parts of the correlator (correlator unit and the antenna unit before the filter) is identical to that assumed in the basic hybrid correlator design. The basic structure for the filterbank processing unit is shown in fig. 1

At full bandwidth, 34 BBCs are used to implement 34 slightly overlapping sub-bands, covering the whole 2 GHz IF bandwidth.

¹If fractional bit shift in the sampler is already implemented for the first generation correlator, this advantage is irrelevant.

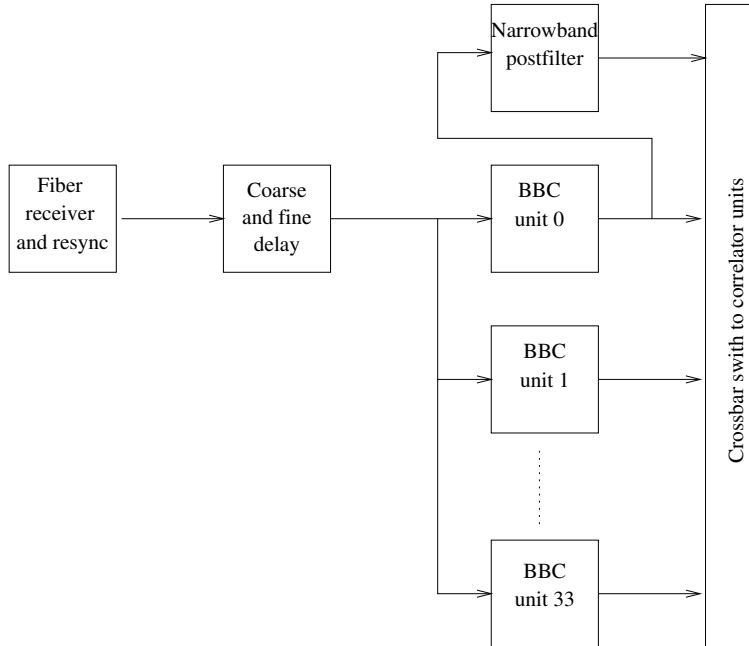


Figure 1: Structure of one filterbank unit of the hybrid correlator. Signal received from the fiber link is compensated for geometric delay, and split into 34 sub-bands. Each sub-band can be freely positioned within the IF. Signal is then transmitted to the correlator units. An extra narrow band filter is used to further narrow one of the sub-bands.

To implement bandwidths narrower than the full IF, less filters are used, and correlator units are cascaded. In this way, bandwidth can be traded with spectral resolution, and the spectral resolution (spectral channel width) varies with the bandwidth. Assuming 1024 channels per IF (32 channels per sub-band), the maximum resolution is 62.5 kHz (1024 channels over a 62.5 MHz bandwidth). The correlator units always operate at full clock rate².

For bandwidth smaller than 1 GHz, it is not possible to efficiently use the 34 units, since they cannot be evenly divided among the sub-bands. Therefore, some hardware is always left unused and the overall usable band is reduced. For example, in table 1 the case of 1/4 bandwidth is considered, and compared with the corresponding configuration in the basic FIR design, using cascaded filters and continue sub-bands.

Parameter	Dig. BBC	Cascaded FIR	unit
Sub bands	8	16	
BW of each SB	62.5	62.5	MHz
Usable BW per sub-band	58.6	62.5	MHz
Channels per SB	128	64	
Usable channels	120	64	
Resolution	0.49	0.49	MHz/ch
Tot. usable channels	960	1024	
Tot. usable bandwidth	469	500	

Table 1: Comparison of different parameters for a 1/4 bandwidth (500 MHz nominal) configuration, using digital BBCs or aliased sub-bands with filter cascading

²The proposed hybrid correlator has 32 channels per sub-band, in full polarization mode (all 4 Stokes parameters computed). For dual polarization mode, the number of channels double. This corresponds resp. to 8192 or 16384 channels per polarization over the 8 GHz bandwidth. In this report, we will always assume the full polarization mode.

We can see that, due to the fact that two sub-bands cannot be used, 1/16 of the total bandwidth is lost. This happens for all the reduced bandwidth configurations, apart from the 1 GHz one (that can be implemented using 17 sub-bands). It is possible to recover part of this loss by cascading together the sharp filters (only one broad filter is needed anyway), in such a way to have less overlapping among the sub-bands. For example, cascading 4 filters one can have a total loss of 1/64 of the total bandwidth (in the case above, an useful bandwidth of 492 MHz or 1008 channels), that is probably acceptable. It is worth noting that, even in this case, only two different filter configurations are required.

To implement bands narrower than a single slice, a single narrowband post-filter can be used. A symmetric filter with coefficient recirculation can implement in principle arbitrary decimation factors. The correlator shift register is then clocked at a corresponding reduced rate, with all correlation resources cascaded together as a single 1024 channels correlator ³.

The number of taps is in fact inversely proportional to the output bandwidth, giving a constant number of multiplications per unit time. The design for such a filter will be subject of a further report. In this way, without channel recirculation, spectral resolution increases linearly with decreasing bandwidth. Apart from the different implementation, these performances are identical to those of the basic hybrid correlator.

To reduce filter size, and cost, a two stage filter has been used. To avoid increasing quantization losses, the second filter uses a many-bit representation for the signal. In the proposed design, 10 bits are used for the signal, and 11 for tap coefficients. The first filter is used to reduce the sampling rate to 1/8 of the input rate (500 MHz), with a passband of 62.5 MHz (the required final passband) and a transition band sufficient to prevent aliasing. In this way it is possible to have a flat passband response and a high stop band rejection with a very limited number of taps (64 in the proposed design). The second filter operates at a much reduced sample rate, and can thus obtain a given performance with a number of taps approximately reduced by a factor of 8. The 128-tap filter proposed here is thus equivalent to a single pass 1024 tap filter. Bandpass shape is determined by the second filter.

Tap coefficients can be determined initially by separate optimization of the two stages, using standard algorithms. Then the second filter is modified to compensate for the small roll-off in the passband due to the first filter.

This approach has the further advantage of increasing the stop band rejection over a significant fraction of the undesired band, thus decreasing noise folding in the passband by ≈ 6 dB.

2.1 Specifications

In this chapter, the main specifications for a digital BBC are given. Each BBC must satisfy the basic specifications for the ALMA hybrid correlator. Some further specifications deal with the capability of frequency tuning. The bandpass specifications derive from the capability of overlapping the sub-bands.

Input data is given as a 16x time multiplexed stream, with an input frequency of 250 MHz, for a total data rate of 4 GS/s (2 GHz bandwidth). Data is represented with 3 bit, using any convenient code. An input format using 32x multiplexed data at 125 MHz is also possible, implementing a 1:2 multiplexing stage internally to the BBC.

Output data is given as a non-multiplexed data stream, at a clock of 125 MHz, using a resolution of 3 or 4 bit, for a bandwidth of 62.5 MHz. Data is rescaled to have maximum efficiency in the correlator, using uniform spaced thresholds.

The output data represents an arbitrary sub-band of 62.5 MHz, SSB converted to baseband, with no frequency folding. The actual bandpass is 30/1024 of the input bandwidth, with guard bands folded around slice edge. For a input bandwidth of 2 GHz, this means that the total bandwidth is 58.594 MHz, (1.953-60.547 MHz), with a guard band extending from -1.95 to +1.95 MHz of each slice boundary. With a 32-channel correlator per sub-band, this corresponds to deleting the first and last channel of each sub-band, keeping the remaining 30 channels. With 34 sub-bands, one obtains 1020 usable channels over 2 GHz, i.e. the first and last 4 MHz of each IF band are not usable.

³Using a Xilinx blockRAM to implement delays and to store tap coefficients, decimation factors of up to 512K can be implemented. This corresponds to a bandwidth of a hundred Hz with a resolution of a fraction of a Hz, well below any conceivable application.

Ripple in bandpass should not exceed ± 0.5 dB. Stop band rejection should be better than 30 dB everywhere, and better than 35 dB on average. This last specification corresponds to the requirement that the excess noise folded in bandpass should not exceed 1% of total noise already present.

Positioning of each band should be possible with an accuracy of 3.6 millihertz (40 bit). In this way fractional bit delay correction should be possible, with a total phase error of less than 1.5 degree over an integration period of 1s. If fractional bit delay correction is not required, the positioning accuracy must allow the correct positioning of the overlap in the sub-bands. In a configuration where four filters are cascaded, the required overlap corresponds to $1/4096$ of the total 2 GHz band, i.e. 0.49 MHz. 13 bit of resolution in the LO are thus sufficient.

All circuitry shall fit in a single XCV1000 chip. For comparison, the interim correlator FIR requires 5 XCV1000 chips for approximately the same performance.

3 Implementation

The internal structure of a digital BBC is shown in fig. 2.

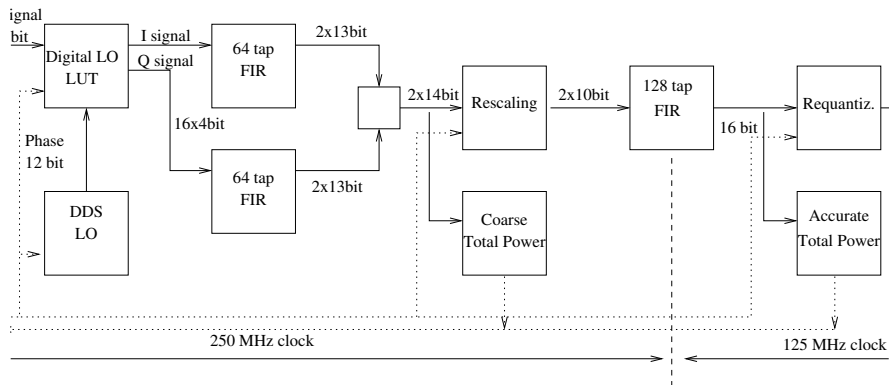


Figure 2: Structure of a digital BBC. The signal is mixed with a quadrature LO, filtered by a first broad filter that implements also phase shift for BBC conversion, re-quantized to 10 bit, filtered by a second sharp filter, rescaled and re-quantized to a final resolution of 3 or 4 bits. Total power meters are used to monitor signal level.

The BBC is composed of a digital oscillator (DDS), a digital quadrature mixer, a first broad band filter and a second sharp band forming filter. The output from the second filter is rescaled and requantized, and the in-band total power is measured.

The mixer selects a region of the input band, that is filtered through the broad band filter. This filter has a bandwidth of 62.5 MHz, equal to the final required bandwidth, but is just sharp enough to allow a decimation by a factor of 8. The output of the filter is thus a bandwidth of 250 MHz, of which 62.5 MHz represents the desired data, and the remaining contains the filter guard bands.

The broad band filter introduces also a phase offset, to reject the unwanted sideband when the I and Q streams are summed together. Its output is applied to a sharp filter, that operates at the decimated sample rate and has a $1/4$ passband. It selects the desired portion of the 250 MHz bandwidth, sharpening the band edges. Both filters have the same final bandwidth, but with different input bandwidth and sharpness.

3.1 Local oscillator and mixer

The local oscillator is a 36 bit DDS, operating at 250 MHz, implemented as a 3-stage pipelined adder of 12 bit per stage. This implies that any frequency change takes effect with a delay of 3 250MHz clock cycles.

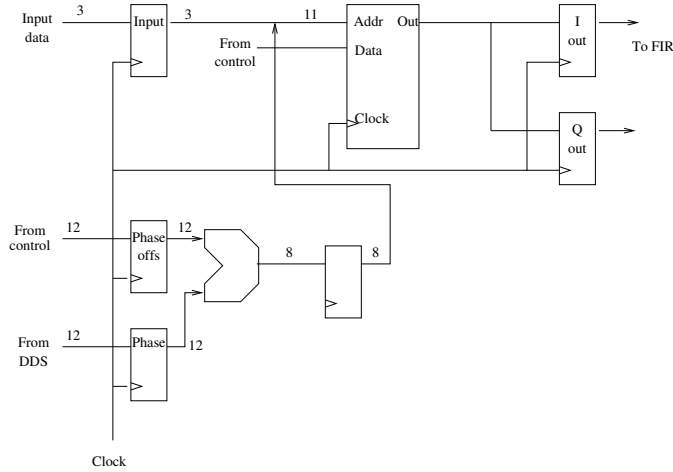


Figure 3: Design of a mixer slice. Each parallel sample is processed in a similar way, with an appropriate value for the phase offset. Multiplication and sine/cosine generation is performed in a LUT memory, loaded at startup.

The mixer slice conceptual schematic is shown in fig. 3. Each parallel sample is independently fed together with 8 bit of phase to a 2048*8 LUT RAM. The RAM output represents the I and Q values of the downconverted sample, with 4 bit of accuracy each. The phase of each slice is offset by a programmable value, to take into account the relative delay with respect to the 250 MHz clock. In this way, extra 4 bits of frequency resolution are available. The DDS and the phase offsets in each slice must be reprogrammed every time the frequency is changed. RAM values are loaded at startup and need not to be changed when the filter is reprogrammed.

If no fractional bit delay is needed, a single stage DDS, with 12 bit resolution, is sufficient. This eases synchronization problems in the programming. The saved complexity is however marginal, and the extra resolution may be useful for other purposes.

RAM values are loaded at startup. Since they represent a table with a fixed representation of sine and cosine values, common to all the 16 RAM, they can be loaded from an external ROM and need not to be changed when the filter is reprogrammed.

3.2 Broad band filter

The I and Q samples are fed to a set of 4 FIR filters, with 64 taps each. The filters compute 2 parallel samples for each of the I and Q signals, representing a x2 multiplexed signal with an effective clock rate of 500 MS/s. Filter conceptual schematic is shown in fig 4.

The filter provides a passband of 1/16 of the input band, with guard bands of 1/8 (folded) on each side. The filter is fixed (the fixed coefficients loaded with the design), since tuning is performed using the digital LO. Avoiding the necessity of dynamically changing tap coefficients the circuit complexity can be significantly reduced.

Each branch provides an additional phase shift of +/- 45 degrees⁴, to implement a digital SSB. The I and Q branches are summed together, to reject the undesired sideband. The filter has a rejection of better than 40 dB over the whole bandwidth, and better than 55 dB over most of the bandwidth. SSB rejection is in excess of 60 dB. Reducing filter coefficient representation to 8 bit (7 bit plus sign) degrades somewhat the performances, but the stop band rejection is always better than 40 dB. Filter response for ideal (infinite resolution) and 8-bit coefficients is shown in fig. 5

For optimal performance, filter bandpass is placed near half the input bandwidth (1 GHz, 1.125-1.1875 GHz in the examples shown here). Band selection is performed tuning the LO from 1.125GHz

⁴Since the filter response is not real, tap coefficients are not symmetric. It is therefore not possible to use filter designs that exploit symmetries.

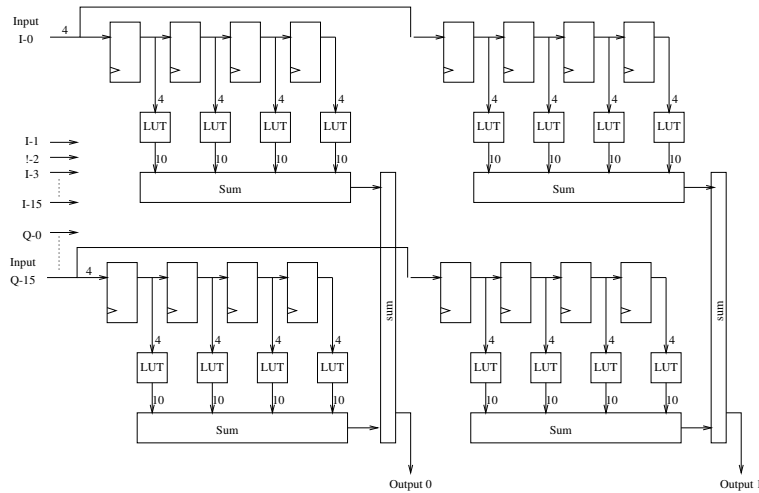


Figure 4: Coarse FIR schematic. Signals from I and Q mixers are multiplied by coefficient taps in LUT tables. Input is from 16 time multiplexed streams, output is to 2 time multiplexed streams.

(band 0) to -0.8125GHz (band 31) (positive and negative frequencies are distinct in a SSB converter).

The filter output is represented as a stream of 10 bit samples (a truncation and justification mechanism is provided to select the 10 more significant bits in any circumstance), 2x time multiplexed, with a data rate of 250 MHz, for a total bandwidth of 250 MHz (but with only 62.5 MHz of useful bandwidth).

3.3 Sharp filter

This stream is then filtered using a 128-tap FIR filter, implemented as two time multiplexed slices with 64 taps each. Coefficient width is 11 bit, and each multiplier produces a 22 bit output. Filter is even, and data from positive and negative lags are summed together to save multipliers. Each multiplier is run at 250 MHz, and computes odd and even taps on alternate clock cycles (coefficient recirculation), to give one output sample at 125 MHz data rate. In this way, only 16 multipliers are used for each slice, and 32 for all the 128 taps. The block schematic for one slice is shown in fig. 7. A detailed description of the multiplexing and decimation algorithm for this filter is given in appendix (chapter 5).

The filter output is 10 bit, at a data rate of 125 MHz. Different tradeoffs are possible for the filter shape. It is possible to obtain an equiripple design with ± 0.7 dB of in-band ripple and an uniform rejection of 30 dB on the unwanted section of the sub-band. With a less equiripple design, it is possible to obtain a better in-band flatness, at the expense of a lower rejection (-27 dB) on a few channels adjacent to the passband. Total band shape (of both coarse and sharp FIRs) for this latter case is shown in fig. 6 left. Sharp FIR provides additional rejection over 3/4 of the total bandwidth, further reducing aliased noise contribution. Shape around the passband is shown enlarged in fig. 6 right, with nominal, guard and passband indicated by ticks above the plot.

Filter output statistics is collected by a digital total power meter, implemented using a hardware multiplier. DC offset is also monitored, and optionally removed.

The filter output is then multiplied by a programmable rescaling factor, and truncated to 3 or 4 bits. In this way, uniform quantization levels with arbitrary spacing can be generated. The rescaling factor is programmed by external logic using total power informations gathered by the total power module.

3.4 Considerations on FPGA resource usage

The XC2V1000 chip can operate at clock frequencies above 300 MHz, but to operate above 250 MHz, several limitations must be observed. At this frequency it is difficult to implement adders longer than 12-14 bit, and block multipliers can handle signals of a maximum of 10-11 bits. This is at the present

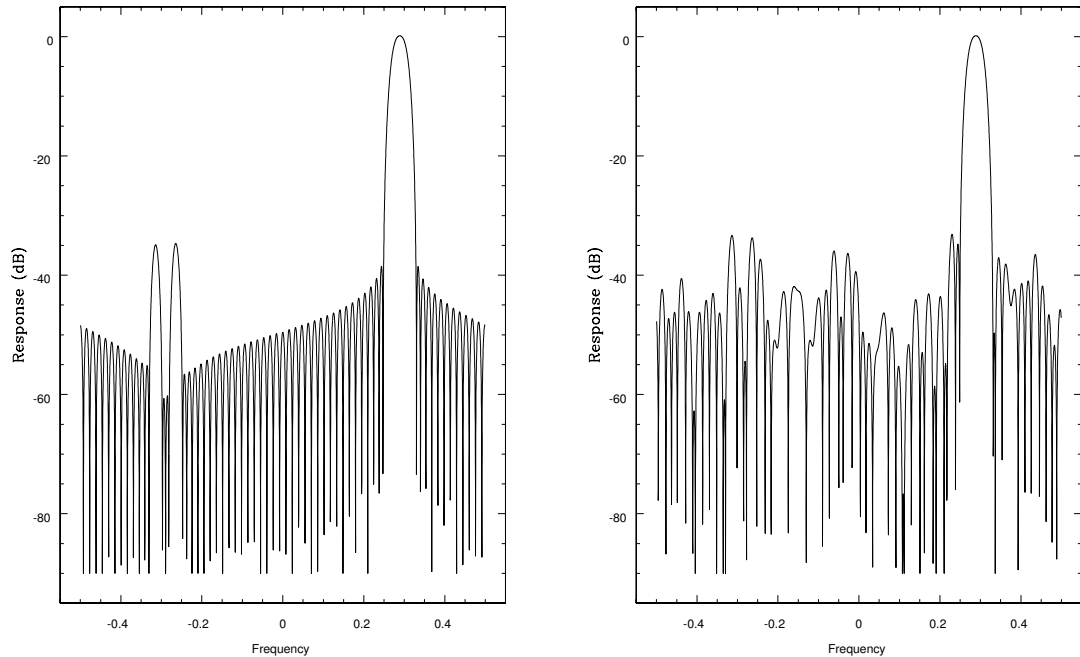


Figure 5: Coarse FIR passband. Left plot is for infinite resolution tap coefficients, right plot for coefficients rounded to 8 bit

the major technological uncertainty of the project. Extended simulation, both using software tools and an hardware demonstrator, is needed to guarantee the technical feasibility of this design.

The most resource intensive part of the circuit is the first FIR. FIR taps are implemented using lookup tables (LUTs), and 10 LUTs are needed for each tap. With 4 FIRs of 64 taps each, this requires 2560 LUTs. Approximately the same number of LUTs are required for the adder chain, totaling 5120 LUTs.

The digital mixer requires one RAMBLOCK and ≈ 45 LUTs for each input stream, and the DDS requires about 150 LUTs. Therefore the LO/mixer requires 870 LUTs.

The second FIR filter is based on block multipliers. A total of 32 multipliers are required (of 40 available), while other 3 multipliers are used by the rescaler and total power circuit. Each section of the second FIR filter requires a total of about 100 LUTs, to implement 4 taps. A total of 32 sections are required (128 taps, 64 times a TMF of 2), for a total of 3200 LUTs.

Control logic, total power and re-quantization circuitry probably does not require more than 200 LUTs.

The grand total is therefore 9400 LUTs, very close to the total number of 10240 available LUTs in the chip. An occupation of 90+ percent is possible, considering the very regular structure of the device, but may reflect on the performance.

If the filter taps need not to be changed (bandwidth selection can be obtained by tuning the LO, narrower bandwidth can be implemented reducing the number of parallel slices processed, and for very narrow bandwidth an extra filter operating at reduced clock rate can be considered), it is possible to reduce the number of LUTs used in the first filter. Half the taps can be represented using only 4 bits, instead of the 8 used in the central section of the FIR kernel. This translates in a saving of approximately 1000 LUTs. Also in the second filter about half the taps can be represented with 4 bits less than in the central core, but in this case the total saving is less than 500 LUTs. Globally, using these savings one obtain a final occupation of 75% of the chip, that guarantees a good routing of the design.

Using larger chips, more than one channel can be implemented in a single FPGA. The larger chips currently available are the XC2V10000, that would host almost 16 channels (at 100% occupation). At the current rate of growth of FPGA size, it appears feasible to host all 34 SSB in a single unit within a

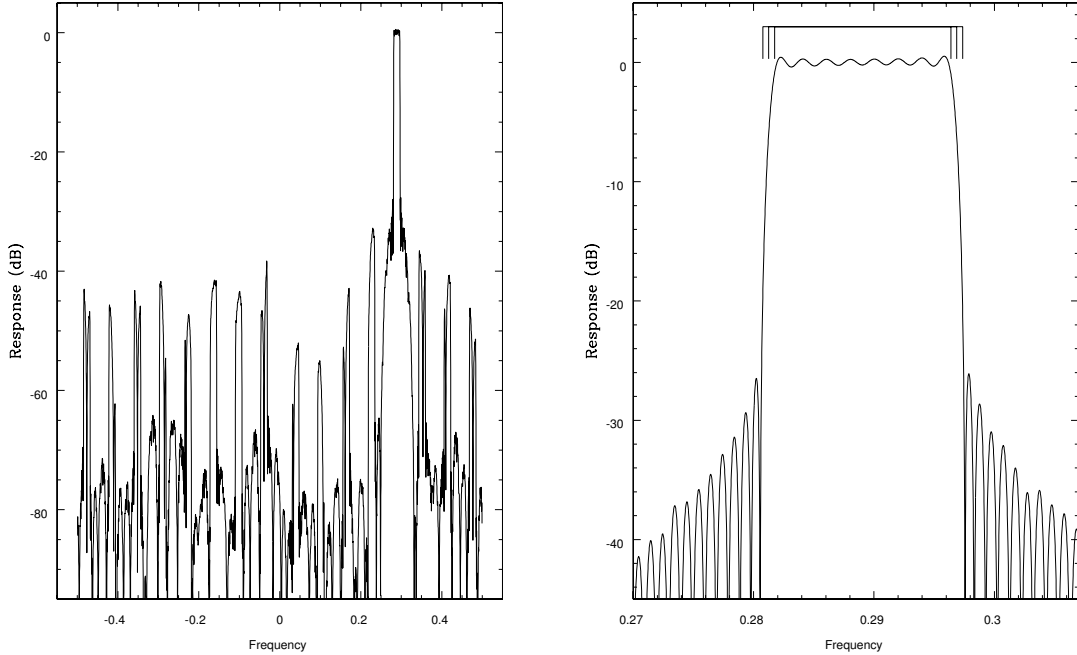


Figure 6: Global filter passband (coarse and sharp FIRs. Left plot is for whole band, right plot is a zoom around the bandpass. Ticks above the band indicate the guard band (wider), nominal (1/32 of IF, intermediate), and pass bands (30/32 of the nominal band)

few years. With the current available devices, 4-5 FPGA would be required for the whole filterbank.

4 Conclusions

A two-stage digital filter with an equivalent number of 1024 taps can be implemented in a single XCV1000 field programmable gate array. A 32-channel digital filter can be reasonably fit into 4 of the larger Xilinx chip available today. A digital LO and SSB converter can also fit in the same chip. Such a design stresses the FPGA capabilities, and requires careful design, but if proved feasible would allow for a simplification and cost reduction of the second generation correlator filter board. A digital LO would allow for band overlapping, and for fractional delay compensation.

5 Appendix: Algorithms to implement time multiplexed symmetric FIR filters

The problem of implementing a decimating symmetric FIR filter is extensively treated in the literature (see for example [10]). When the filter is time multiplexed, however, new problems arise.

A filter computes the sum $s_j = \sum w_i r_{j-i}$, with, in the symmetric case, the tap coefficients satisfy the relation $w_i = w_{-i}$. For a decimating filter with a decimation factor D it is sufficient to compute s_j for $j = kD$. Therefore, the sum can be decomposed in

$$s[kD] = \sum_m \sum_{n=0}^{D-1} w[Dm+n] r[D(k-m)-n]$$

where the second sum can be performed using a single multiplier/adder that cycles through a set of D

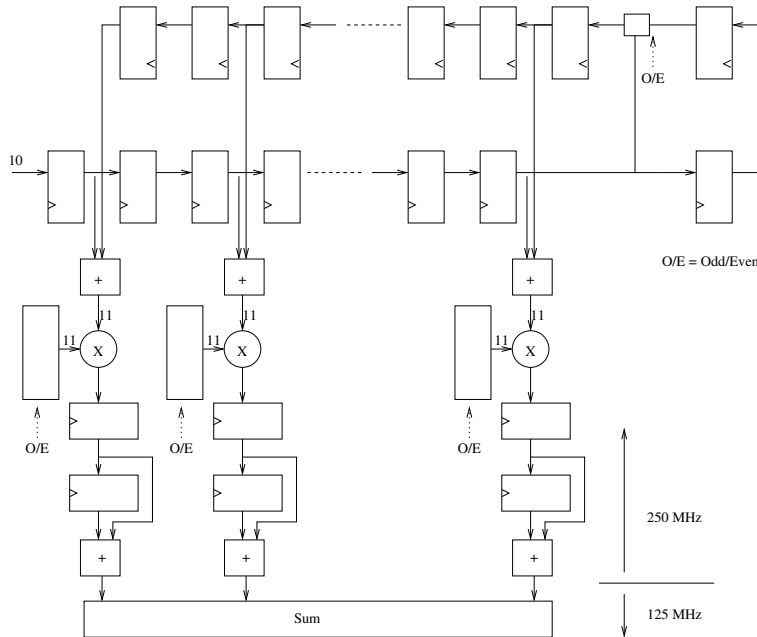


Figure 7: Schematic of the second (sharp) FIR. Only one time multiplexed slice is shown. Delay line is folded back to exploit filter symmetry. O/E is a signal to distinguish between odd and even clock cycles.

consecutive tap coefficients. In this way, a multiplication is performed at each undecimated clock cycle, and the total number of taps is equal to D times the number of multipliers.

If the filter is asymmetric, no complication arises even when the filter is time multiplexed, i.e. the input signal consists of alternate samples fed to parallel FIR filters (filter branches), whose outputs are added together. Each branch processes a subset of the required taps, and the total length and decimation factor is the corresponding value for each branch multiplied by the number of parallel branches.

If the filter is symmetric, and the symmetry is used to reduced the number of multipliers, the situation is more complex. In this architecture, the delay line is folded back, and corresponding samples are summed together before being multiplied by the tap coefficients. Samples must be reversed in *Last In/First Out* (LIFO) order to present them to the multipliers while the correct tap coefficients w_j are applied.

If the filter is time-multiplexed, the situation is still more complex, since the delay line folding is different for different branches, and the folded section may come from a different branch.

In this chapter, we will examine the particular case of $D = 4$ and $\times 2$ time-multiplexed symmetric filter, i.e. the sharp filter described in this design. In this case, the LIFO has a depth of 2, and can be implemented with two delay cells and a switch.

The topology is completely different in the two cases were the filter has an odd number of taps (zero-delay tap w_0 , and $(N - 1)/2$ positive and negative taps), or an even number of taps ($N/2$ taps placed symmetrically around the zero delay).

The simplest case is that of an odd number of taps. In this case, the even branch and odd branch are completely independent.

The sequence of operations for a 15 tap filter is shown in figure 8 and tab. 2. In the table, the filter taps are listed on top, and the corresponding sample indexes used to compute the output value s_8 are listed below the corresponding coefficient. The computation is done in two branches (odd and even) and two consecutive clock cycles (T_1 and T_2). On each cycle, samples are advanced through the delay line, and each multiplier/adder computes one term of the filter sum. After two clock cycles, all samples are advanced by two positions, corresponding to four index values, and a new product is started.

The total number of multipliers required is four, two on each branch. The products shown in the outermost sections of table 2 are computed by the first multiplier in the two branches, while the second multiplier computes those in the central section.

Weights	w_7	w_6	w_5	w_4	w_3	w_2	w_1	w_0	w_1	w_2	w_3	w_4	w_5	w_6	w_7
Cycle T_1															
even branch				12				8				4			
odd branch			13				9		7				3		
Cycle T_2															
even branch		14				10				6				2	
odd branch	15				11						5				1

Table 2: Sequence of products for a $2\times$ time multiplexed symmetric FIR filter with an odd number of taps. Sum for output sample s_8 is computed in four steps (two filter branches and two clock cycles).

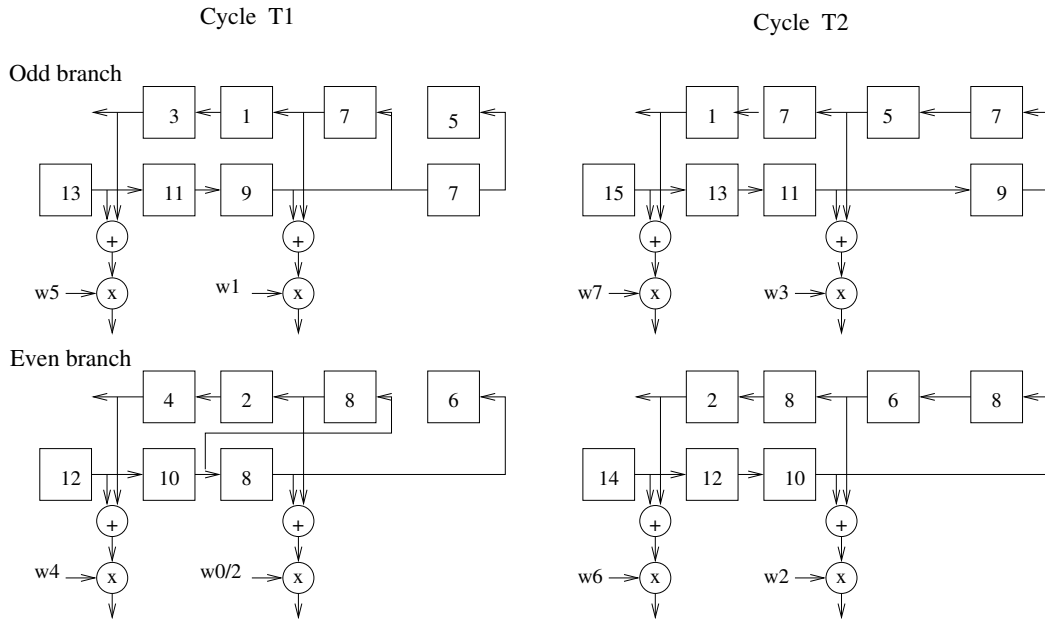


Figure 8: Hardware architecture for the filter described in table 2. The beginning of the backward delay line (positive taps) is alternatively connected to the end of the forward delay line (negative taps) or to an extra 2-cell delay line, to implement a 2-stage LIFO register.

The backward delay line (positive taps) is connected to the forward delay line (negative taps) through a LIFO implemented as a two-cell delay line and a switch. In this way, samples in the backward delay line are interchanged in pairs, and compare in the right order at the multiplier input. Corresponding samples (e.g. sample 3 and 13) are summed together before being multiplied by the corresponding tap coefficients. Products for the two cycles are accumulated together, and the result for all multipliers and branches are summed together after T_2 . The two branches are slightly different. In particular, the even branch has one less delay cell before the backward line, to properly align the odd and even samples. Sample corresponding to coefficient w_0 is summed twice. Therefore the corresponding coefficient is halved.

For an even number of taps, the two branches are no longer independent. The coefficients w_i and w_{-i} , numerically equal, apply to samples in different branches of the time-multiplexed filter. This means that the same tap coefficient w_i applies to samples of opposite parity for negative and positive taps, as can be seen in table 3.

In the hardware implementation, this means that the two branches must have a cross connection, as seen in fig. 9. The forward delay line of each branch is fed to the backward line of the *other* branch. This may complicate the design, since long connections imply long propagation delays.

Weights	w_8	w_7	w_6	w_5	w_4	w_3	w_2	w_1	w_1	w_2	w_3	w_4	w_5	w_6	w_7	w_8
Cycle T_1																
odd branch				13				9	8				4			
even branch			14				10			7				3		
Cycle T_2																
odd branch		15					11				6				2	
even branch	16				12							5				1

Table 3: Sequence of products for a $2\times$ time multiplexed symmetric FIR filter with an even number of taps.

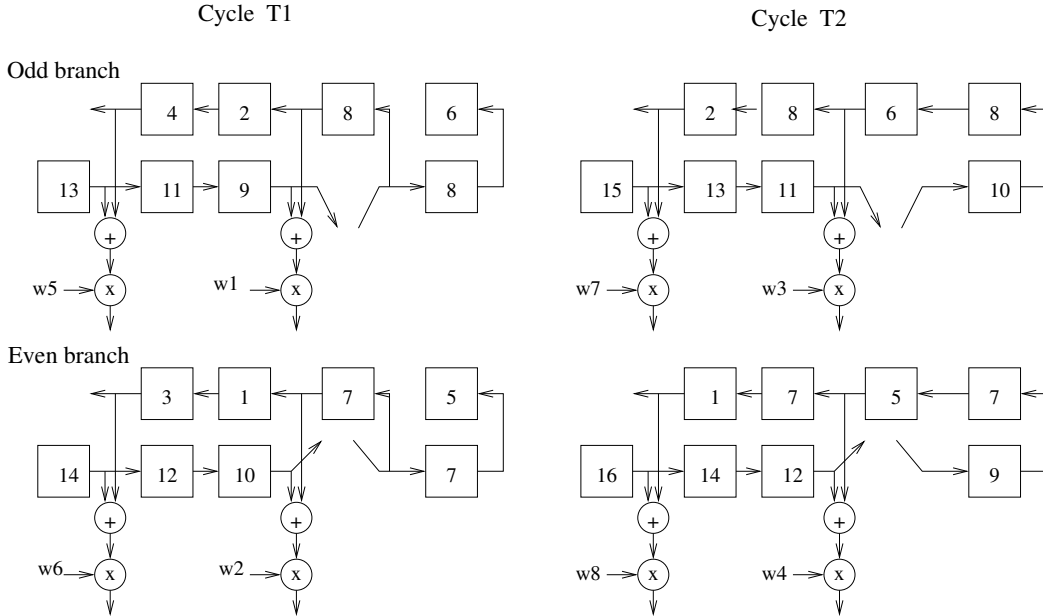


Figure 9: Hardware architecture for the filter described in table 3. The backward delay line of each branch is connected to the forward delay line of the *other* branch.

References

- [1] A. Baudry, A.W. Gunst: "ALMA Filter Bank Specifications and Delay Tracking", ASTRON report (in preparation) (2001)
- [2] B.R. Carlson, P.E. Dewdney: "Efficient wideband digital correlation", Electronics Letter, IEEE, **36-11**, 987 (2000)
- [3] B.R. Carlson: "A Closer Look at 2-Stage Digital Filtering in the Proposed WIDAR Correlator for the EVLA", NRC-EVLA Memo 03 (2000)
- [4] B.R. Carlson: "Refined EVLA WIDAR Correlator Architecture", NRC-EVLA Memo 14 (2001).
- [5] B.R. Carlson: "WIDAR Correlator Sensitivity Losses", NRC-EVLA memo 26 (2001)
- [6] G. Comoretto: "A digital BBC for the Alma interferometer", Alma report n. 305 (1999)
- [7] G. Comoretto: "Possible designs for a hybrid correlator", Arcetri Internal Report n. 8/2000
- [8] R. Escoffier and J. Pisano: "Test Report of the Baseline ALMA Correlator Digital Filter" Alma Report n. 409

[9] B. Quertier: "Proposal for a Future Correlator Filter Board", ASTRON report (in preparation) (2001)

[10] Harris HSP43168 FIR filter data sheet

Contents

1	Introduction	1
2	Architecture	1
2.1	Specifications	3
3	Implementation	4
3.1	Local oscillator and mixer	4
3.2	Broad band filter	5
3.3	Sharp filter	6
3.4	Considerations on FPGA resource usage	6
4	Conclusions	8
5	Appendix: Algorithms to implement time multiplexed symmetric FIR filters	8