New Technologies in VLBI
ASP Conference Series, Vol. 306, 2003
Y. C. Minh

Development of a Digital Base Band Converter (DBBC): Basic Elements and Preliminary Results

Gino Tuccari

Istituto di Radioastronomia CNR, Contrada Renna, 96017 Noto (Sr), Italy

Abstract. VLBI acquisition terminals will undergo an important improvement, due to the fact that disk based recorders have just been introduced. This fact has great impact on the efficiency and reliability of the entire system, building upon the increased performance already achieved with the MKIV formatter. In this renewal process, it seems valuable to concentrate attention on the base-band or video-converter, the section of the VLBI terminal devoted to down converting portions of the IF receiver band to base band frequencies. New technologies can be used to help in the process to redesign this part of the system, and a fully digital solution may be considered to be an attractive and realistic viable implementation. This paper illustrates some of the common methods adopted to perform this kind of functionality, taking into account 'digital radio' equipment, which operates using similar principles. The required performance is relevant with respect to commercial available devices and this imposes the adoption of ad hoc solutions for implementing a good part of the functionality.

1. Introduction

An optimal approach to record digital data from a radioastronomical receiver should not take into account the prior division into small sub-bands, as current VLBI terminals do. Such channelization is due to the limited amount of bandwidth recordable with a magnetic tape. Thus, in the past, to widen the observed band meant an increase in the bandwidth of a single channel, and especially to increase the number of recorded tracks. Such an approach, even if still valid, does not optimally fit with a now possible and entirely re-thought system. We have learnt, from the story of VLBI in these past decades, that the upgrade process of this technology must be gradual, and not traumatic, in order to guarantee compatibility with existing systems and correlators.

We think this lesson is still valid, and the conversion to a fully digital terminal must pass through the small multi-band method, before changing in approach to use a wide band recording and processing system. This passage will probably arrive with the advent of the new generation of correlators. In this process, the redesign of the frequency conversion with digital methods should be taken into account, as fast commercial samplers are now available with appropriate clock frequency and quantization levels, and recent FPGA products allows

sharing of the fast digital conversion and filtering operations needed to perform the necessary functionality. Indeed, a full set of base band converters, such as those placed in the current MKIV/VLBA acquisition terminal, may be seen very schematically as a chain composed of a fast sampler, a fast connection bus, a set of individual conversion/filtering sections (the order of these operations can be inverted or mixed), an interface to an existing formatter or formatting unit – which act as as a standard interface, a controller, and a communication section for setting purposes.

During the last few years, IRA has developed in Noto a prototype providing some of the functionality required of a fully operational set of video converters. The main reason for such a development derived from a project where radio astronomy stations without any VLBI terminals were involved; thus, there was a necessity to overcome the lack of any existing conversion facility. While this project, and the related technological research, has produced what is necessary to satisfy the original goal, it is still unsatisfactory for the different set of required functions to be performed with a traditional VLBI terminal – but it could represent a good starting point in a dedicated project. A description of this prototype is given here, with some details about its realisation, and some preliminary results.

Several partner institutions have shown interest in supporting our activity to study and develop a system to emulate an entire set of analog video converters, with the double advantage of replacing what can be in some way considered an obsolete component in existing terminals, and of giving the option to those stations without any VLBI equipment, to access this observing methodology with a relative low cost solution: a digital base band converter set plus a disk based recorder. Thus, the time is ripe, from both points of view – technological and of opportunity – for the decision to proceed with an executive project that could drive the VLBI community to convert their terminals to a fully digital implementation in a few years, with great promise. This conversion would potentially benefit, along with additional advantages, those stations interested in converting their receivers' signal managing from the analog to the digital domain, which would then reduce the cross-talk, and frequency and amplitude dependence over the long distances between the antenna and control room.

2. Analog Base-Band-Converters

The MKIV video-converter (VC) and the VLBA base band converter (BBC), as they are commonly called to distinguish between them, both convert a window from the IF receiver output down to the video band, or in other terms, to a frequency range covering the region from almost zero up to the maximum bandwidth of 16 MHz. This conversion is a frequency translation, using the phasing method of single side band conversion. Both upper and lower sidebands are shifted, preserving the information, and they are then low pass filtered.

In the case of the MKIV VC (NEROC 1978), it has been designed to operate with an IF frequency range of 100–500 MHz, and a video output range from about 800 Hz up to $8-4-2-1-0.5-0.25-0.125-0.0625$ MHz. In this case, the converter consists of an image rejection filter followed by a video amplifier, low-pass passive filters and again another amplifier. Seven-pole Butterworth

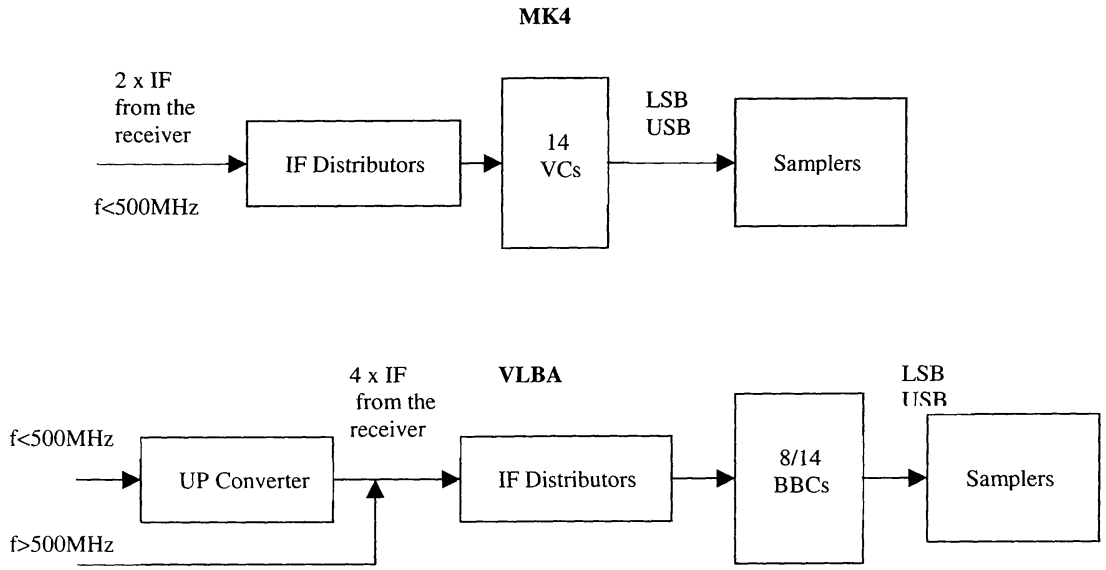

Figure 1. Base Band Converter in the MKIV and VLBA environments.

low-pass filters are used, being considered to be about optimal for VLBI, due to the relatively rapid cut-off, with a very smooth phase response. The smooth, almost linear, phase response ensures a close phase match between different video converters. The total loss percentage, with an optimum location of the 3 dB point (in percent of the bandwidth) of 91%, is 3%. The −3 dB cut-off frequency is placed at 7.2, 3.6, 1.8, 0.9, 0.45, 0.225, 0.112, and 0.0562 MHz. At 1.2 times the cut-off frequency, the response is down more than 10 dB. The in-band ripple is less than 0.5% peak-to-peak over a large part of the band. In the case of a VLBA BBC (NEROC 1987), the input IF frequency is in the range 480–1020 MHz, so that for receiver frequencies in a lower range, an up-converter is adopted. Even in this case, the frequency conversion is performed in a single step, still using the phasing method for obtaining single side band reception. Active filters are adopted, with an 8-pole Butterworth typology. The bandwidth is selected in the range 62.5 kHz up to 16 MHz. The overall performance is similar to what we have seen in the MKIV VC.

In Fig. 1 a schematic view is shown of the collocation of the converters inside the MKIV and VLBA environments.

3. The Digital Down Converter

A typical direct digital down converter accepts an input signal, sampled at a high rate, and then down converts the desired frequency band of interest to base band. Moreover, the sample rate is adjusted by a factor, taking into account the final bandwidth. Digital down converters are fundamental components of digital radios, as they perform the critical frequency translation needed to shift the required information to a lower range, where it can be detected or recorded.

A digital down converter consists of an oscillator, which is a Numerically Controlled Oscillator (NCO) or a Direct Digital Synthesizer (DDS), a complex mixer, and a low-pass filter. Many authors have reported on this and a com-

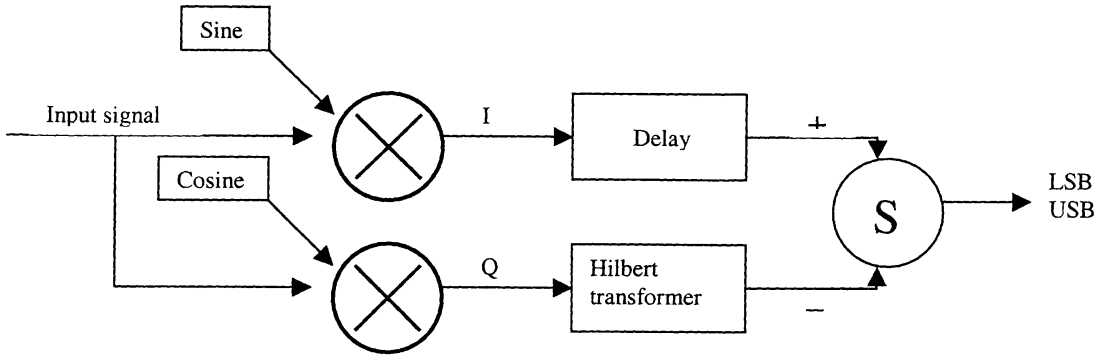

Figure 2. Complex mixer for the SSB phasing method converter.

pact summary may be found in the Xilinx documentation (Xilinx 2002). The band-limited output from the filter allows the reduction of the sample rate by decimating. In terms of system performance, a critical component in the digital conversion is the Numerically Controlled Oscillator. This component generates a sampled digital sinusoid, that is mixed with the incoming signals to shift the frequency spectrum. We would like to recall the difference between this method and a traditional frequency conversion realized with a real mixer. The latter uses a single mixer unit in such a way that signals at the frequency sum and difference, between the input signal and the local oscillator, are produced. Upper and lower sidebands are converted to the same baseband, and then superimposed. When a single sideband is required, a complex mixer can be adopted, or a double single mixer with two components in quadrature (Fig. 2); in this case, two converted quadrature components are obtained. This conversion is called the phasing method, and differs the advantage of producing a simple shift of the frequency band to the desired range, in an amount equal to the frequency carrier of the local oscillator.

In order to get the upper and lower side band channels, the separated components (in phase, and in quadrature) are summed and subtracted. In the quadrature section, a Hilbert transformer is added, while in the 'in phase' arm, a delay section is inserted to compensate for what is added in the other branch. The Hilbert transform is such that it shifts all positive components by +90 degrees with respect to the local oscillator, and all negative components by −90 degrees. This method is similar to what is used in the VC and BBC analog implementations. Insufficient precision, or inaccuracy in the sinusoid, produces a poor signal-to-noise ratio product and spurious spectral artefacts. Thus, careful attention has to be applied in the quantization representation. The frequency synthesizer is an accumulator used to summate phase increment values. Such a method is relevant, because the synthesized signal frequency is not in an integer relationship with the clock frequency, which is typically equal to the sample rate. Arbitrary frequency values can then be obtained without modifying the sampler rate. Moreover, the phase increment and initial phase setting allows the proper local oscillator to be generated with absolute phase control.

The NCO module generates a digital approximation to a sine and cosine wave, with a shape of a staircase. To be realized, a phase accumulator and a

sine/cosine look-up table (LUT) are used. A second register stores the increment value that is added to the phase register at each clock cycle. The integrated phase value, or a truncated version of it, is used to address the sine/cosine LUT, whose output represents the amplitude of the generated signal, used as the local oscillator. The phase can be expressed as the time integral of the angular frequency ω, such that a change $\delta\theta = \omega\delta\tau$, where $\delta\tau$ is the time interval over which the phase change takes place. It is then possible to express the output frequency of the sinusoid to generate, in terms of the phase change and the clock rate of the phase accumulator:

$$f = \frac{\delta\theta f_{clk}}{2\pi}.$$

In the NCO implementation, the maximum possible phase value is governed by the phase accumulator, that can contain at maximum the number 2^n, where n is the number of bit representation. Then, a full representation of the range $0, 2\pi$ is to be done in the range $0, 2^n - 1$. The phase increment must be included in the range $0 \leq \Phi \leq 2^n - 1$, and finally the generated frequency is:

$$f = \frac{\Phi f_{clk}}{2^n}.$$

It can be seen that the word length of the phase change sets an upper limit on the achievable output frequency, while the accumulator word length sets the frequency resolution of the oscillator. Other considerations are related to the noise generated by the oscillator. The main sources for such noise are due to the amplitude and phase quantization. The effect on the amplitude is reflected on the SNR, which is related to the word length by about $SNRdB = 6n + 1.8$; while the phase quantization produces spurious artefacts, and the Spurious Free Dynamic Range (SFDR), which is the ratio between the power of the carrier and the strongest spurious product, is $SFDRdB = 6n$.

The mixed signal has to be filtered to isolate the portion of spectrum containing the needed signal. A typical filter is a narrow band one, relative to the input unselected band, with a high rejection of the unwanted spectrum. It then looks convenient to perform such preliminary filtering by adopting a decimation scheme. A commonly used approach is to insert a Cascaded Integrator-Comb (CIC) filter, that is not other than a recursive implementation of a moving average filter, or 'boxcar.' Its spectral response is the *sinc* function ($sinx/x$). In such a filter, the number of effective taps is an integer multiple of the decimation ratio, and the filter nulls alias into the pass-band when the spectrum is folded by decimation. If the pass-band is narrow enough the rejection of the aliased image is quite good. Several sections can be cascaded to lower the amplitude of the side lobes. The pass-band of this filter presents a relevant roll-off, that has to be corrected by a clean-up filter. A narrow pass-band filter in the final section improves the alias rejection, and the roll-off compensation. An important issue is that the filter response, referred to the output sample rate, is almost independent of the decimation ratio, and so a single clean-up filter can be adopted for different decimation ratios. Identical filters can be adopted for both the in phase I and the in quadrature Q mixer outputs.

A more complete description can be consuldted in Fig. 3, where it can be seen that the sample rate can be adjusted to match the video frequency band,

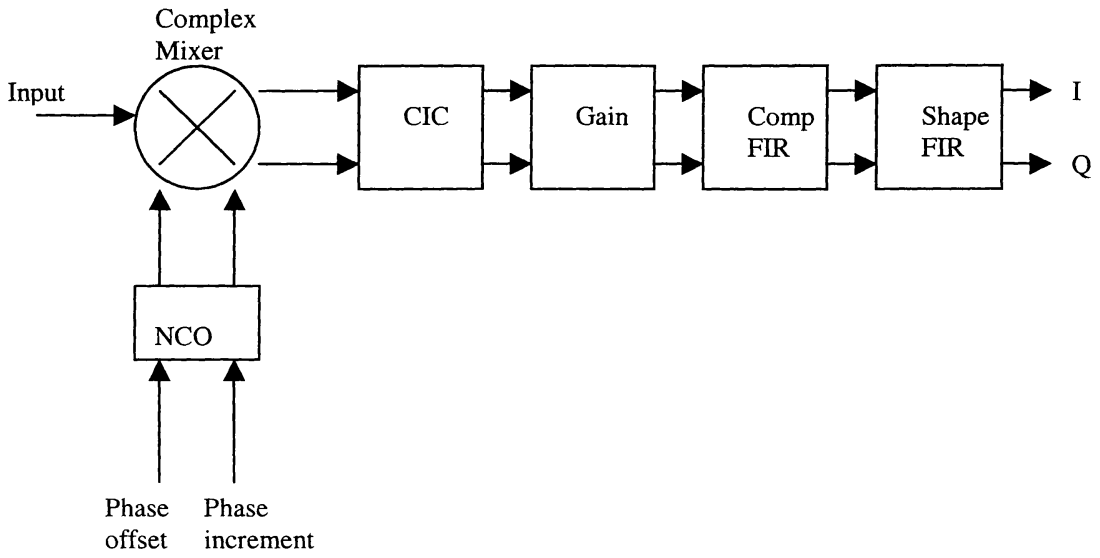

Figure 3. Digital Down Converter, a typical implementation.

by adopting a multi-stage, multi-rate filter consisting of a cascade of CIC filters and two polyphase decimators. A coarse gain filter is added for gain adjustment, a compensation filter to correct the CIC bandwidth, and a programmable filter for selecting a variable piece of band.

With the Numerically Controlled Oscillator, the f_{out} is:

$$f_{out} = \frac{f_s \delta\theta}{2^{B_\theta}},$$

and the corresponding phase increment is:

$$\delta\theta = \frac{f_{out} 2^{B_\theta}}{f_s},$$

with $\delta\theta$ as the phase increment, f_{out} the frequency generated, f_s the sampling frequency, and B_θ the number of bits used in the phase accumulation register. The resolution δf is defined by:

$$\delta f = \frac{f_s}{2^{B_\theta}}.$$

The CIC filter, or Hogenauer filter (Hogenauer 1981), is a multi-rate element used to realize a large sample rate change in digital systems. The important feature is that they are multiplier-less structures, consisting only of adders, subtractors, and registers. Their use is efficient in the design, when the sample rate is much larger than the bandwidth occupied by the converted signal. In Fig. 4, the structure of a CIC filter is shown. The integration part consists of N integrator stages operating at the high sampling rate f_s, and each stage is implemented as a one-pole filter with unity feedback coefficient. The transfer function is:

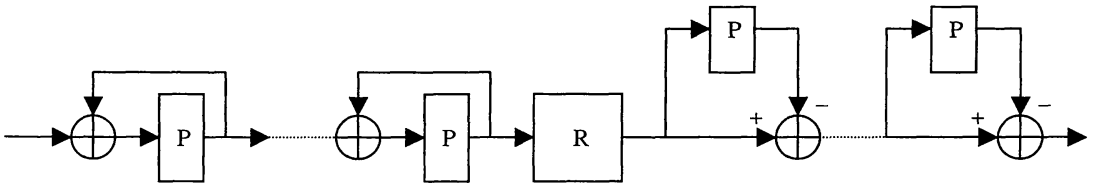

Figure 4. Cascaded Integrator-Comb Filter, a generic implementation.

$$HI(z) = \frac{1}{1 - z^{-1}}.$$

The comb section operates at low sampling rate f_s/R, where R is the integer rate change factor. This section consists of N comb stages, with a differential delay of M samples per stage. This differential delay is used to control the filter frequency response. The transfer function, referenced to the sample rate f_s, is:

$$HC(z) = 1 - z^{-RM}.$$

A rate change switch is included in between these two sections, and is nothing other than a sub-sampler of the output of the last portion of the integrator stages, reducing the sample rate from f_s to $\frac{f_s}{R}$. The overall transfer function is:

$$H(z) = HI^N(z)HC^N(z) = (\sum z^{-k})^N.$$

It may be noticed that the CIC filter is then equivalent to a cascade of N uniform Finite Impulse Response (FIR) filter stages with unit coefficients, or, in other terms, is equivalent to a cascade of N boxcar filters. The low-pass characteristics of the CIC element are evaluated by considering:

$$z = e^{\frac{j2\pi f}{R}},$$

evaluated at the low sampling frequency $\frac{f}{R}$. The magnitude response of this filter is:

$$|H(f)| = (\frac{\sin \pi M f}{\sin \pi \frac{f}{R}})^N.$$

A compensation filter can be used to adjust the bandwidth response. Such a filter should give a $(x/sin(x))^N$ response.

4. A VLBI Implementation

What we have described is a common architectural view implemented in numerous digital down converters today. Due to the fact that these configurations are

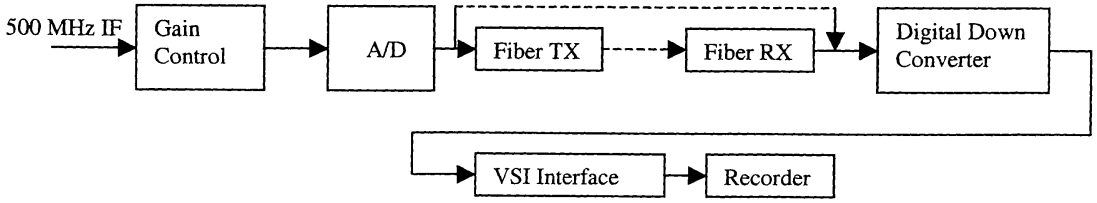

Figure 5. A possible digital conversion for VLBI.

very often performed using Field Programmable Gate Array (FPGA) structures, and that the hardware is being implemented through software solutions, a radio system with a similar design is often called Software Radio. More specifically, for radio astronomy purposes, several projects are related to the development of large international programs. All these have specific requirements strictly related to their particular applications and they constitute an interesting and similar effort that the radio astronomy community is performing in this area, with different approaches (e.g. Comoretto 2002). Let's evaluate what would be a similar application in a VLBI terminal in use today. Roughly, we need a digital multi-channel system with single side band separation, with both USB and LSB availability, operating with an input band about 500 MHz wide. Such a converter must shift slices of the incoming band with good performance, in terms of spurious free dynamic range, and separation between adjacent channels. And, still very important, it has to take care of the delay and phase control.

Consider Fig. 5, which is a schematic view of a system able, in principle, to perform the functionality of a video converter. A 500 MHz band is presented at the input side of an A/D conversion stage. The first part of such a stage is represented by an adjustable gain amplifier, in order to have appropriate levels at the inputs of the actual converters. Following this conversion, which is executed at a sample rate satisfying the Nyquist requirements, samples B0 bits wide are made available in parallel format, on a high-speed bus. Multiplexing in two or more blocks, a reduced clock frequency could be applied.

Whether or not this A/D stage is to be placed in the receiver area, the transmission of the sampled data could proceed through a fast serial fibre transceiver. This solution could be particularly useful in those stations where the control room with the acquisition equipment is far from the receiver position. In such cases, long distance connections through cables result in a frequency-dependant amplitude, overall amplitude reduction, and any other effects that this could cause. A digital transmission can resolve, or minimize, such effects.

After the reconstruction of the signals, or else just after the sampler, the data bus enters a fast I/O module. From now on, the description concerns a single channel, because the other elements operate in parallel; exactly as now, in the analog base band converter, the modules operate in parallel. The sampled data, along with the clock, are mixed with a local oscillator signal, before proceeding to the filter stage where the data are re-sampled. Further actions are to adjust the shape and the gain. Every operation described here has to be real-

ized in the complex domain, because the phasing method has to be applied in order to shift in frequency the desired portion of the band, while discriminating between the upper and lower sides. For this purpose, a Hilbert transform should be considered in the quadrature chain, while a delay stage with proper length should be the equivalent element in the 'in phase' branch. The filters with appropriate bandwidths should be applied to exactly define the shape of the final band. The last part should be to gather the different channels, with appropriate 1 or 2-bit data representation to form a unique collection of bit streams under a standard interface protocol, the most convenient, being VSI, before sending the data to the transfer channel, or to the recorder.

5. A Prototype for VLBI Observations

A first prototype has been realized, and is still under development, but some preliminary evaluations may be presented. The developed system is part of a project related to a near-real time radar VLBI acquisition terminal. In this project, several ex-Soviet radio telescopes observe echo signals and need to record just from the receiver IF output. Taking into consideration a band of 500 MHz width, the goal was to down convert a single portion of narrow band from any point of the IF spectrum, with a step of 0.0625 MHz and a bandwidth selected from 2 MHz down to 62.5 kHz, in octaves. This system has been successfully proved in a single station up to about 300 MHz input, with an output filter of 2 MHz, and with LO resolutions of 1 MHz. The higher part of band, and the smaller resolutions, are currently under advanced development.

This acquisition terminal may be thought of as a coherent system, because the signal is converted and filtered using a high-stability phase reference, recorded on disk, and the data files are transferred through the network time-tagged in order to correlate them at a common place. In other words, it is a simple narrow-band VLBI terminal, because the echo signals from the radar transmission are monochromatic, and the limited bandwidth is in this particular situation good enough to transfer the data on a standard Internet network, without the big bottleneck due to the transit time.

In Fig. 6 is shown a schematic view of the system, that reflects what we already have seen in the other sections of this paper. Of course, more details are introduced here, since they refer to an actual prototype.

The output of the receiver must present a constant power level, and a well-filtered anti-aliasing filter. Some gain adjustment will be realized in the chain, but, for a proper working system, these requirements are mandatory. The first stage is a commercial Maxim A/D evaluation board able to sample at a maximum rate of 1.5 GHz, producing a digital version with 7–8 bit data width. In our prototype, three-four bits are used, and the sampling rate has a fixed value of 1.024 GHz, produced by a synthesizer phase-locked to an external 10 MHz reference; the amplitude is fixed at +4 dBm, and the clock signal is fed to the converter module through a narrow band filter to reduce the higher order spectral contribution. Particular care must be taken with regard to the phase noise contribution of such a clock generator; SSB noise components exceeding values of about −90 dBc/Hz can degrade the overall conversion quality. The sampled data are then presented at the output of the module, de-multiplexed at

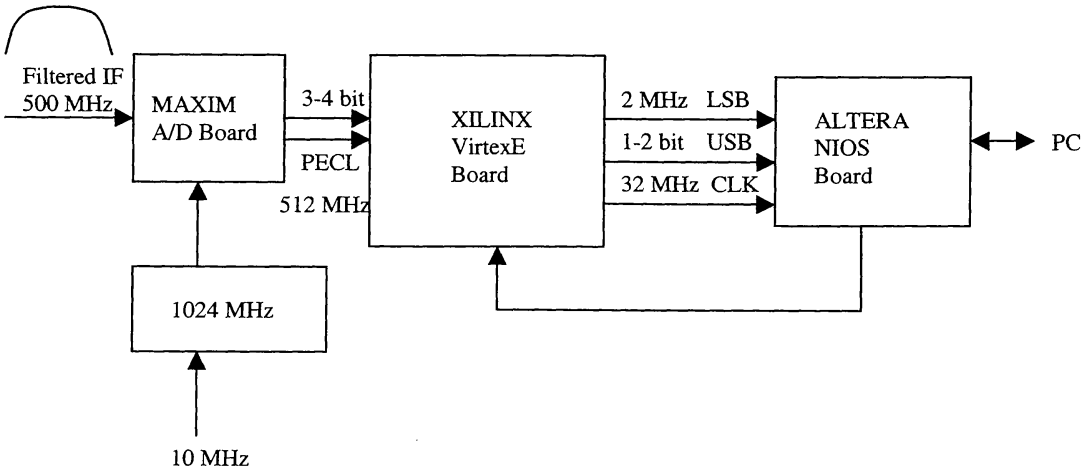

Figure 6. The Noto DBBC prototype, schematic view.

half data rate with output PECL levels. Such a differential standard is one of the possible I/O methods adopted by the Xilinx Virtex-E family devices, which are used in this prototype as high-speed interface logic, and as a general purpose Field Programmable Gate Array element.

Data transfer to the Xilinx chip is assured by a dedicated differential fast channel, with a maximum data rate of 622 Mb/s. Each of the four sampled bits is transferred using two lines, along with the reduced clock signals. Manual floor-planning has been adopted to configure the front-end stages of the FPGA chip to meet the required very high performance. Fortunately, only a few elements are required to work at high frequency, and a reduction is quickly applied by distributing the logic in parallel blocks. This functionality is the most critical in the project (Chinnery et al. 2001). In Fig. 7, a schematic view of the internal data processing is shown.

A further de-multiplexing by four is performed, with the clock frequency at 64 MHz, enough to process the Xilinx routing in automatic fashion. Groups of sixteen temporally adjacent data samples are then presented for each of the three or four bits of the representation, as a parallel block of information. Next, the complex mixer is fed by two streams: the sampled signal, and the local oscillator signal produced by the Numerically Controlled Oscillator. Further considerations will be given later about the realization of the local oscillator. Each element of the mixer then processes in phase, and in quadrature, the product between the four-bit representation of the input data and the phase value of the local oscillator. Their product is scaled down to produce a four bit representation, which is then sent to the parallel integration section of the combined integrator cascade filter. This part is fairly unusual when executed in parallel, and it represents a particular implementation of the serial integration process, as described in the previous sections. The realization is done with a polynomial calculation performed with an LUT. In Fig. 8, it is schematically shown how this section operates. At its output, the signal is re-clocked at the reduced sample rate of 32 MHz, realizing the functionality R described. The

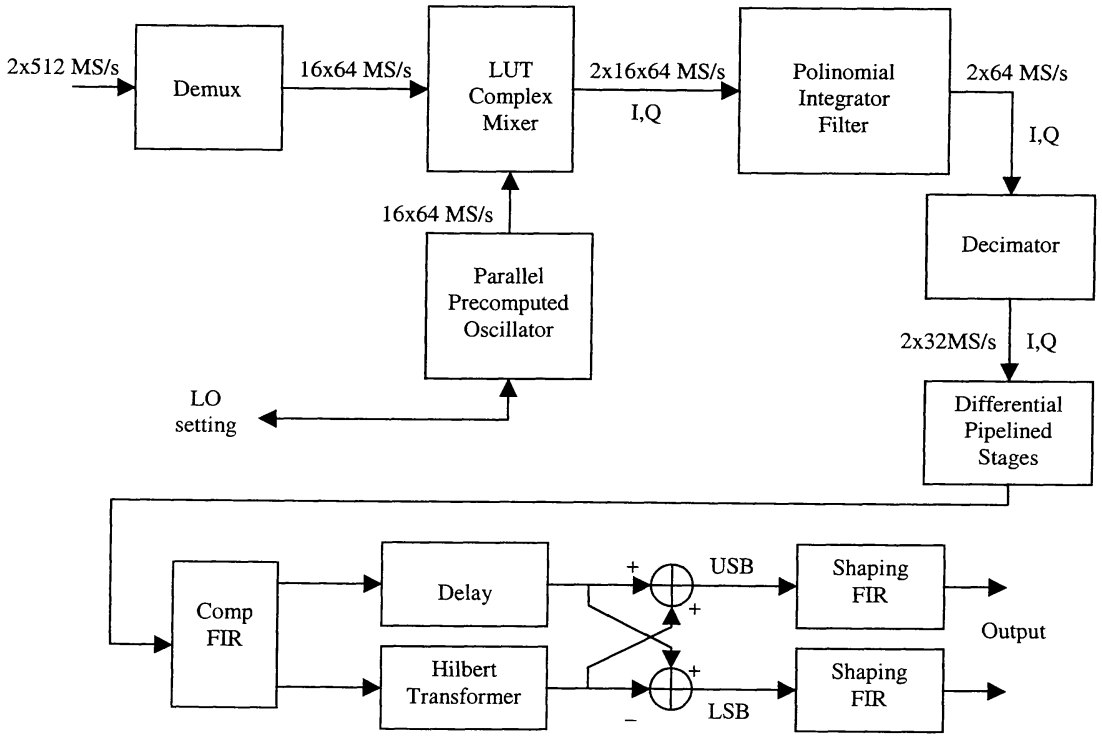

Figure 7. The internal processing of the FPGA data flow implementation.

following section is 16 blocks of differential stages, pipelined to complete the CIC filter implementation. The 32 MHz clock frequency is used for any of the possible bandwidths.

A compensation filter is added at the output of each branch of the complex chain, in order to adjust the shape of the CIC response. Such a filter is also used to rescale the overall gain of the system. The configuration adopted for each FIR filter is the so-called 'inverse' (Chapman 1996), and it is displayed in Fig. 9. One can see that it is a fully pipelined version of the most traditional implementation. This configuration is particularly useful for a numerous taps version, where the long adders can seriously affect the speed performance. We decided to adopt such an architecture everywhere in the design, for the simplicity and homogeneity it presents in the routing process.

Next in our description, we must evaluate the Hilbert transform filter added in the quadrature section of this down converter. It is represented by a different version of a FIR filter, with an appropriate coefficient. It may be useful to recall that this filter is able to shift the phase of the incoming signal by $\pi/2$. This implies that the delay in this path is constant with respect to frequency, and that this additional delay is compensated for the 'in phase' part of the converter with a delay unit able to realign the two streams. Those are then summed and subtracted, one with the other, to form the upper and lower side components. A final 64-tap long FIR filter is adopted to produce a shaped version of the final bandwidth. Even in this case, the implementation is done by adopting the

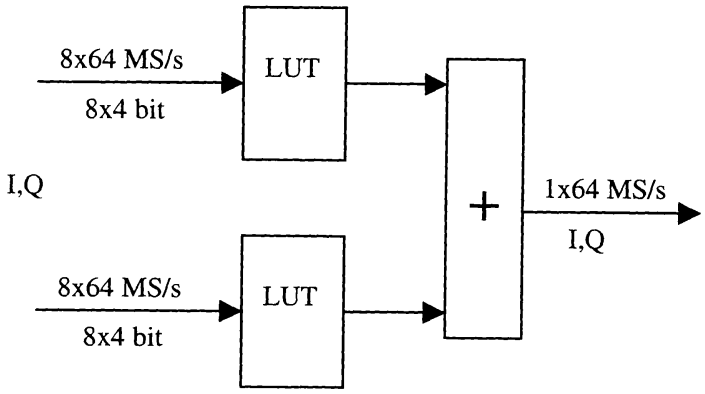

Figure 8. The polynomial integrator implementation.

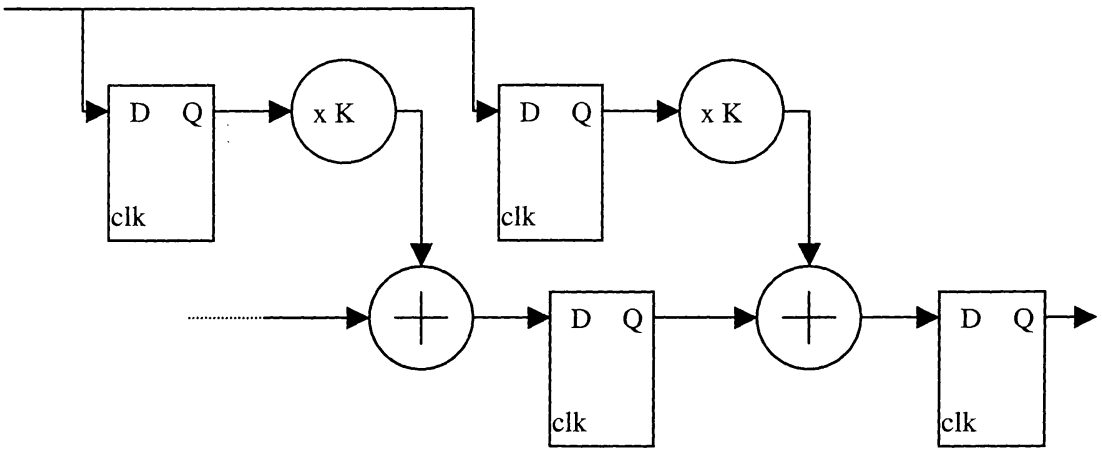

Figure 9. FIR inverse filter.

inverse structure FIR filter. The outermost part of the long filter is truncated at the desired word precision to one or two bits.

The local oscillator implementation using traditional methods seems to be much too problematic, and some modifications have been introduced in order to overcome these difficulties. Indeed, a common NCO should operate at the sampling frequency to accumulate the phase variations. A solution has been found by applying a method that could be called: Parallel Pre-computed Oscillator. In such an implementation, the more traditional direct digital synthesis is performed by parallel blocks of 16 integrators each, loading contiguous cycles in time, at 16 multiples of the phase increment. The initial phase, externally pre-computed, is set for the 16 time slots at the initial stage, when a local oscillator frequency is set, along with the common phase increment. Such computation is then performed at a reduced rate.

In Fig. 10, a picture of the prototype is shown. The leftmost board placed on top is the A/D unit, then under this, in sequence, are the FPGA board, and

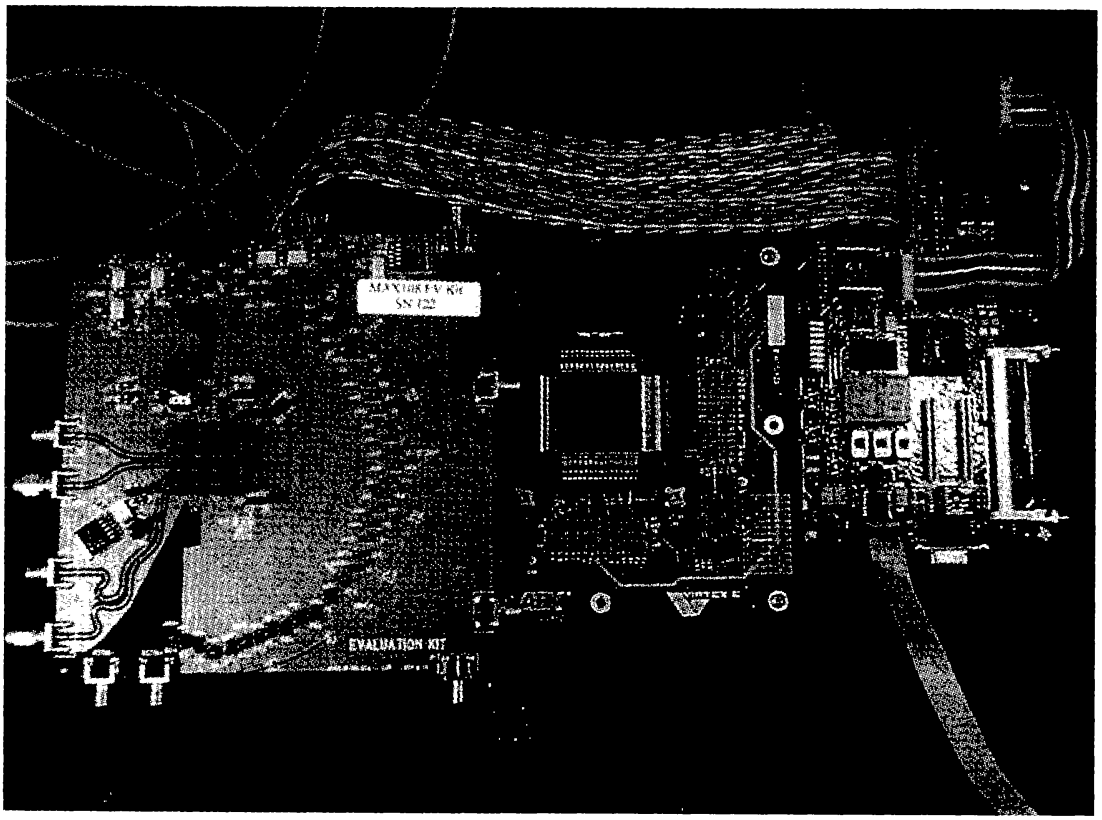

Figure 10. A picture of the prototype.

an Altera Nios processor board for the local oscillator settings and communication. It is important to point out that this system is mostly software, and thus band changes are done through reloading the FPGA configuration, rather than reloading the filters' coefficients. This solution is convenient, because an optimisation comports to modify even the hardware implementation. In Fig. 11, are shown an example of the 2 MHz converted band, with receiver noise in the input; and a line from a satellite echo observed during the July 2002 Radar VLBI observation at Noto. The digital output has been re-sampled and recorded on disk, and with further off-line software, FFT-processed to produce the power spectrum of the output signal. The spectral resolution in the FFT computation is 2 Hz. A direct phase measurement has not yet been done, as we plan to insert a D/A converter in the final stage, to measure the phase transfer function.

6. A Full BBC Replacement: How It Could Appear

A complete set of videoconverters, having the entire set of possibilities now present on the existing analog version, could be developed taking into account the prototype described here. A better approach might be to reconsider the project, by having in mind the realization of a multichannel system, because some parts could be common. An important improvement may be achieved by taking into consideration the new Xilinx family, VirtexIIPRo, for its superior performance. We can suppose that a high-speed board with one or two chips can include the entire functionality of a 14 BBCs down conversion, with the addition

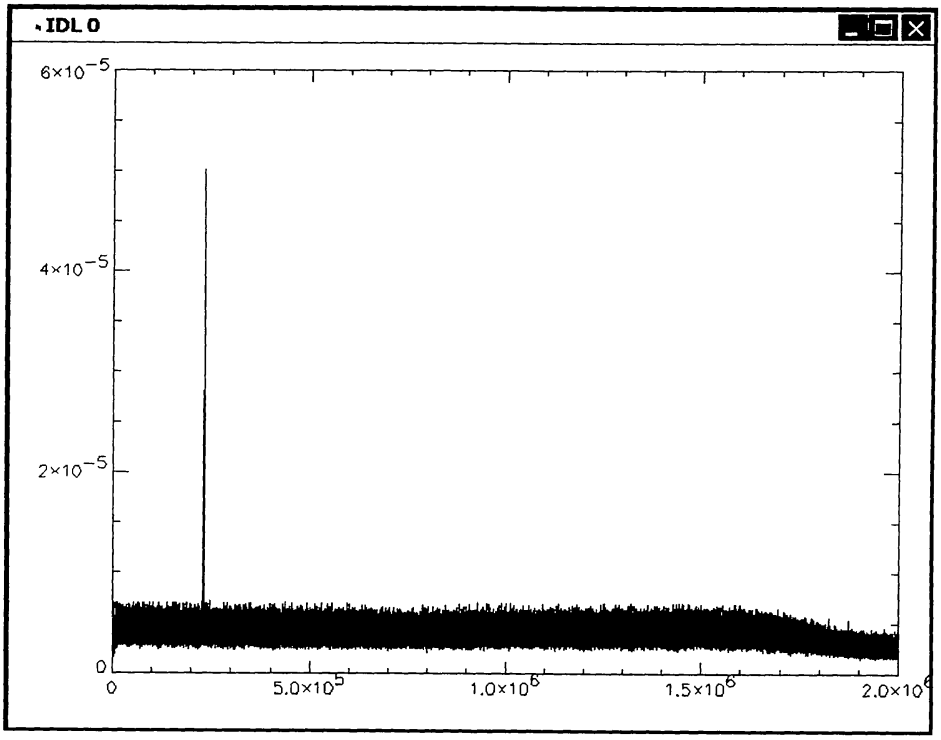

Figure 11. A radar line observed with the prototype.

of an A/D module, a serial link connection, and communication components. The project could take advantage of the new FPGA family for its embedded processor capability, which is useful for the local oscillator synthesis and the communication management. A bandwidth between 50 and 750 MHz could be considered to be divided between multiple band inputs. Thus, for retaining compatibility with present systems, we need to wait for an entirely unique, high band, real-time VLBI data acquisition and correlation process. This same hardware could be updated to perform a few high-bandwidth channels, instead of the multiple low-bandwidth channels version that has been described here. The update would simply consist of an FPGAs software reconfiguration.

Acknowledgments. I would like to thank those who showed interest in this project, and encouraged its development. This includes the EVN Board of Directors, and, in particular, the following VLBI institutions: MPI in Bonn, JIVE in Dwingeloo, NRAL in Jodrell Bank, and IRA in Italy. Particular thanks to Franco Mantovani for supporting us since the beginning of tise idea, Arno Freihold and Gianni Comoretto for interesting discussions, and Paul Burgess for his suggestions, and for revising the text.

References

Chapman, K. 1996, Xilinx Application Notes

Chinnery, D., Nikolic, B., Keutzer, K. 2001, in Proceedings of the 37th Conference on Design Automation, 637

Comoretto, G. 2002, Arcetri Technical Report No. 5
Hogenauer, E. B. 1981, in IEEE Trans. Acoust. Speech Signal Processing 29, 155
NEROC 1978, MKIII Hardware Manual
——— 1987, VLBA Acquisition Rack Manual
Xilinx Inc. 2002, Virtex Handbook

Discussion

N. Kawaguchi: Why do you use the BBC and LPF? The same processing is possible by simply BPF processing. What is your advantage? Tunability? Frequency tunability we can do at the stage of the first SKY/IF frequency conversion.

D. Ferris, G. Tuccari: The two digital filter systems described were designed and built in the early to mid-1990s, when the only suitable technology available worked at a maximum clock rate of 32 MHz. The first example, X-FILT, had an input bandwidth of 16 MHz, and a data rate of 32 Msps. This matched the technology, so we did indeed use simple BPF technology. The more advanced LBA DAS has a 64 MHz bandwidth, and a 128 Msps input data rate. By choosing a particular quadrature downconverter structure (BBC), we were able to divide the input data into four independently processed data streams of 32 Msps, which were then compatible with the technology. It also enabled us to use multi-rate processing, which greatly reduces the total computation required to produce high quality filters.

J. Romney: All three speakers have emphasized the advantages of dispensing with a complex and expensive part of the current systems: the baseband converters. I'd like to ask all three to comment on where the new limits will lie on the number of independently tunable "channels," i.e., sampled frequency bands. Is it as sample as just the number of computations required to filter the full set of channels, or are there other considerations?

D. Ferris, G. Tuccari: In a critically sampled FIR digital filter or filterbank, the number (rate) of computations (multiplications) required to produce a given passband shape depends only on the input bandwidth, while the amount of temporary storage, i.e. memory, is proportional to the bandwidth reduction factor, M. If an FFT is involved, the computation and storage for that part grows with $M\,logM$. In current COTS technology, these factors together mean that 1 GHz of bandwidth may be processed into up to 4 thousand channels in two FPGAs, without adding external memory. An additional consideration is that the number of bits per sample must grow with $\sqrt{M}$ in order to maintain the input dynamic range.

A. Roshi: The question is related to your statement that no bandpass calibration is needed with a Digital BBC. Does this mean that the front-end receiver system will not contribute to the bandpass change?

D. Ferris, G. Tuccari: In general, both phase and amplitude ripples, and instabilities, scale inversely with bandwidth. By defining the narrowest passbands with digital filters, these effects are largely eliminated. Of course, the receiver

front-ends, in fact all parts of the signal path, will make some contribution, but since their bandwidth is typically two orders of magnitude larger, the effects are correspondingly small. The experience of our spectral-line astronomers is that it is necessary to do only one passband calibration per session, rather than once for each source. In this way, they are able to observe twice as many sources in a given amount of time.