



Institutet för rymdfysik  
Swedish Institute of Space Physics

# **EISCAT\_3D DELIVERABLE 9.2: SIGNAL PROCESSING SUBSYSTEM REPORT**

Gudmund Wannberg  
Swedish Institute of Space Physics  
Box 812  
SE-98128 Kiruna, Sweden  
[ugw@irf.se](mailto:ugw@irf.se)

Walter Puccio  
Swedish Institute of Space Physics  
Box 537  
SE-751 21 Uppsala, Sweden  
[walter.puccio@irfu.se](mailto:walter.puccio@irfu.se)

2009-06-05

## Table of contents

<b>1. Introduction.....</b>	<b>3</b>
<b>2. Receiver system signal processing.....</b>	<b>4</b>
<b>2.1. Overview of signal flow.....</b>	<b>4</b>
<b>2.2. Time domain beam-forming/beam-pointing; FPGA beam-formers....</b>	<b>5</b>
<b>2.3. Band-limiting and down-sampling; digital down-converters.....</b>	<b>13</b>
<b>2.4. Coherency processors / triggers for interferometry.....</b>	<b>15</b>
<b>2.5. Quick-look analysis.....</b>	<b>16</b>
<b>2.6. Polarisation matching.....</b>	<b>17</b>
<b>3. Transmitter system signal processing.....</b>	<b>18</b>
<b>3.1. The basic transmitter module.....</b>	<b>18</b>
<b>3.2. RF waveform generation requirements .....</b>	<b>19</b>
<b>3.3. Beam steering and beam-forming.....</b>	<b>20</b>
<b>3.4. An example of a commercial DUC – the AD9957.....</b>	<b>21</b>
<b>4. References.....</b>	<b>-25</b>

## 1. Introduction

In the early days of incoherent scatter, the processing of the radar receiver output voltage into range-gated spectral parameters, and the storage of those at a physically acceptable time resolution, were always major system bottlenecks. When the first EISCAT systems were designed in the 1970s, megawatt-class klystrons, big reflector antennas and parametric amplifiers were readily available, but high-speed A/D conversion technology was still in its infancy, computers were very expensive, painfully slow and totally incapable of handling the receiver data directly, and hard-disk storage technology was only being introduced. The only feasible way to handle the signal processing problem was to restrict the receiver pass-band to the maximum extent possible without losing part of the ion-line spectrum, digitise the band-limited signal, process and time-average the resulting (50 – 100) ksamples/s data stream in application-specific hardware and finally store the Correlator output on nine-track magnetic tape at an effective recording rate of a few tens of kilobytes per second. In fact, achieving the computational throughput required to do full justice to the transmitter and receiver capabilities was still very much non-trivial even when the EISCAT Svalbard Radar was being designed in the early 1990s.

However, the enormous progress in magnetic data storage technology during the last decade has now made it possible to record raw data on affordable off-the-shelf hard-disk arrays at rates of hundreds of megabytes per second. Also, advances in digital integrated circuit technology have changed the personal computing scene completely. It is now possible for the average scientist to have an almost incredible amount of data processing power sitting on her office desk – her PC or Mac is probably fitted with at least 1 Gbyte of RAM and running at a dual-core CPU clocking at 2 GHz or better. It is also likely to sport in a Terabyte or more hard disk storage. If that is still not enough, the local university data centre can probably offer fast, on-line-accessible disk storage in the hundreds of terabytes or more, capable of recording streaming data at rates exceeding several gigabytes per second...

An important consequence of these developments is that EISCAT users now have ready access to tools that allow them to analyse, in a relatively short time, vast quantities of raw data to which little or no pre-processing has been applied. This has triggered a demand for full-bandwidth amplitude-domain data from the new radar, and research into the development of analysis algorithms operating directly on these.

At the same time, the continuing expansion of mobile telephony services worldwide and the ongoing development of new mobile systems standards have spawned a dramatic growth in the mixed-signal integrated circuit arena. A large number of programmable digital receiver chips and up-converter/DAC integrated subsystems, primarily designed for use in mobile systems base stations, are now available off-the-shelf at very moderate cost. These devices are typically designed to interface directly to a control processor and are ideally suited to fill the gap between the analogue RF domain and the digital, discrete-time baseband domain.

If the job at hand is such that a lot of signal processing power is required over and above that built into the mixed-signal chips, Field-programmable gate array (FPGA) technology offers a flexible solution. FPGA circuits with more than 1000 multiply-accumulate (MAC) cores clocking at 550 MHz are now available; a single chip can deliver more than  $4 \cdot 10^{11}$  MAC operations per second.

This Deliverable presents the outline of a possible signal processing architecture for the EISCAT\_3D radar system, including both the part associated with the receiver and the part related to the transmitter waveform generation and beam-forming. While other solutions are of course possible, the particular one proposed here capitalises on the ready availability of inexpensive mixed-signal circuits and FPGA technology, while at the same time retaining full flexibility and programmability in all relevant respects, providing band-limited data as input to several low-bandwidth system functions, and addressing users' wishes to get data from the system in as unprocessed a format as possible for special applications.

## **2. Receiver system signal processing**

The receiver and signal processing systems developed in the course of the EISCAT\_3D study have been designed to:

- digitise the signal voltages generated by the individual array elements as close to the antennas as possible, using a sampling scheme providing at least 30 MHz of non-aliased instantaneous bandwidth,
- subject the data to only the absolute minimum of processing required to enable its storage, and
- leave the application-specific extraction of physical parameters from the data to the end user.

More information on the receiver front-end design can be found in Reference 1. As will be shown below, storing all the data produced by the individual array elements separately is presently a practical impossibility (and probably not meaningful anyway), but by applying a few steps of pre-processing, the data rates can be brought down to manageable levels without losing essential information in the process.

### **2.1 Overview of the signal flow in the receiver system**

Figure 1 shows a simplified view of the planned signal flow through the signal processing system, from the output of the individual ADC to the access point into an optical backbone network (disregarding for the moment the buffering of data on hard media). The abbreviations SG 1...SG m refer to the m first 49-element sub-arrays, each of which generates  $(49 + 49)$  80-Ms/s data streams that are delivered to the signal processing system.

Three clearly distinguishable pre-processing functional blocks can be identified in the flow diagram:

- 1) Time-domain beam-forming / beam-pointing,
- 2) Band-limiting and down-sampling,
- 3) Coherency detection and interferometry triggering

Each of these stages will now be considered. An actual FPGA-implementation of a beam-former will be described in some detail, while the other stages will be discussed more at the conceptual level, to demonstrate that the hardware needed to build a full specification 3D signal-processing system is available today.

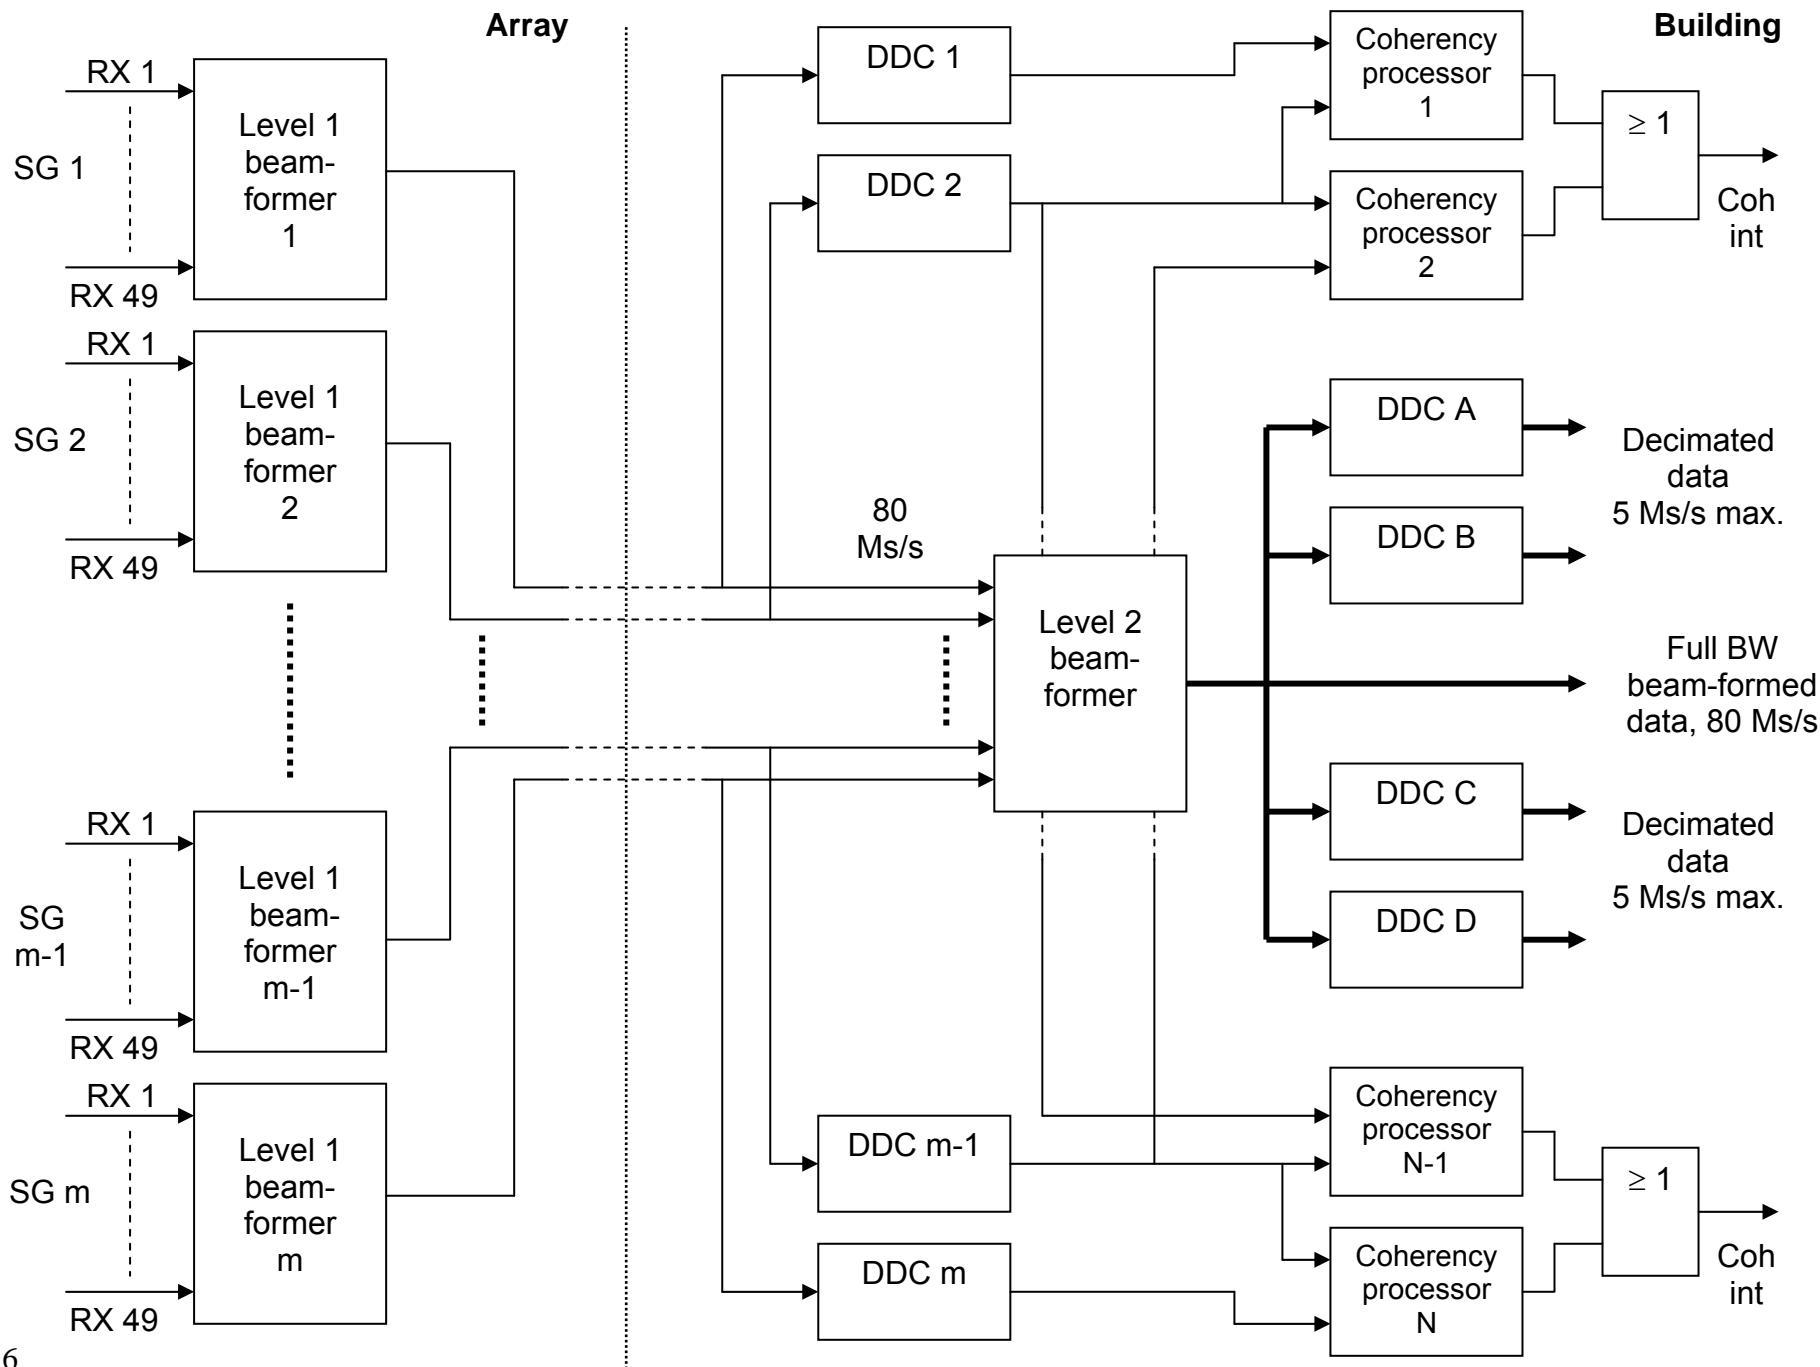
Substantial progress in FPGA technology is to be expected before the design of the full 3D system is finalised. The final signal-processing system architecture is therefore likely to end up looking rather different from the one proposed here, probably a lot more compact – but it would still have to provide at least the basic processing functions listed above!

## **2.2 Time domain beam-forming / beam-pointing; FPGA beam-formers**


The fully populated Core array will contain about 16000 X-Yagi element antennas. Each receive-only array will contain at least half this number, such that the total system will comprise about 48000 elements! Each element will be equipped with a dual channel (one channel for each of the two polarisations) direct-sampling receiver system, covering a bandwidth of at least 30 MHz and generating two continuous 16-bit sample streams. *Constructive undersampling* (aka *sub-sampling*) will be employed, with the signal band falling in the sixth Nyquist zone. The exact sampling rate can therefore be fixed only after the frequency allocations for the 3D system have been finalised, but will fall in the 80 – 90 MHz range.

The total sample rate generated in the Core is thus in the order of  $2.9 \cdot 10^{12} \text{ s}^{-1}$  – a totally ridiculous number, even given today's data storage facilities! One way to transform this into a more tractable problem is to process the receiver outputs in a *beam-former*, where the data streams from the  $\approx 16000$  individual elements are combined into a very few (at most 6-10) composite data streams, each running at the same 80-90 MHz rate. Each composite data stream includes contributions from all array elements and represents a fully focussed beam pointing in some desired, freely selectable direction.

Apart from making the system data rate manageable, the described *beam-forming* process is also appropriate from an information-theory point of view; in a radar system, meaningful signals will only be received from those regions in space that are illuminated by the radar transmitter and so only data representing beams that look into those regions need to be stored.



**Figure 1:** Schematic diagram of a subset of the data flow through the EISCAT\_3D signal processing system, from the individual ADCs to the storage/network units.

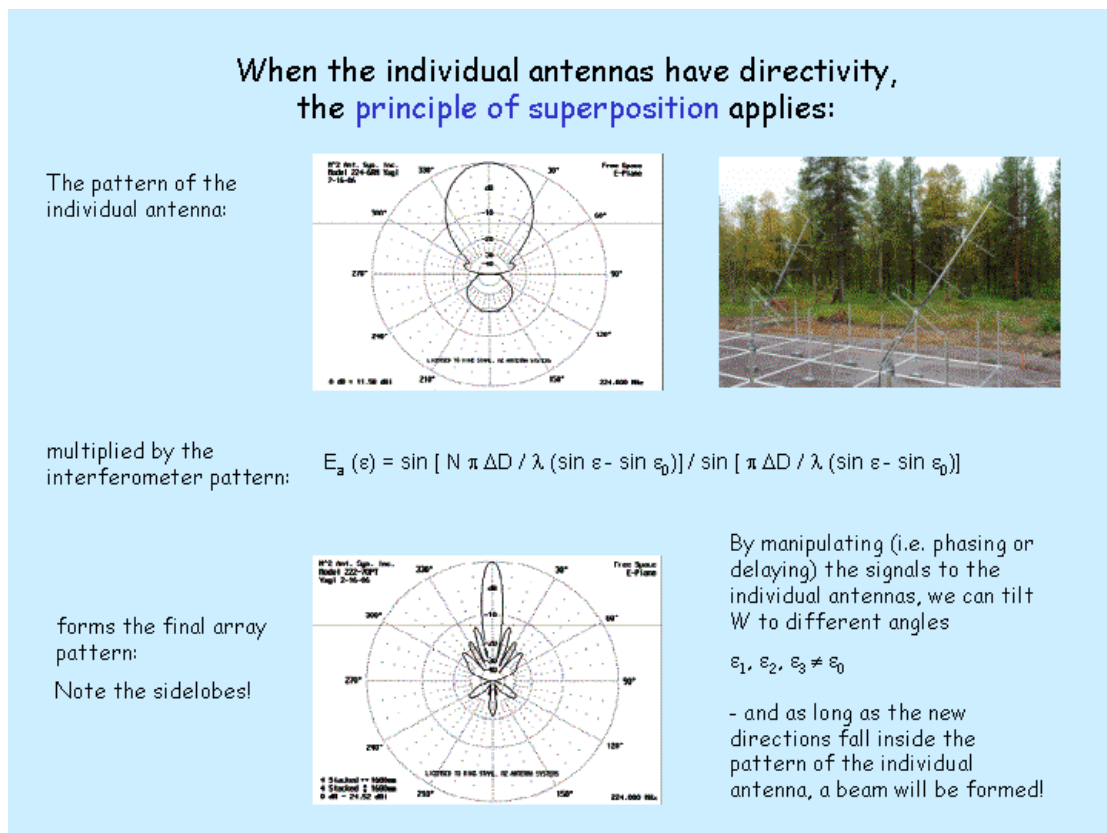


**Figure 2:** Illustrating the phased-array principle.

Steering the beam is conceptually very simple. The signals from the individual antennas are manipulated (i.e. phased modulo- $2\pi$  or time-delayed) such that only those components of the individual signals representing a wave front entering the array from some desired direction add constructively. Signal components associated with waves arriving from other directions are then added more or less destructively and suppressed; in one dimension the end result is an *interferometer pattern* described by the equation in Figure 3. If the inter-element distances are properly chosen with regard to the wave front tilt angle relative to the array plane, only a single (lowest-order) interferometer lobe is formed. In the specific case of the EI\_3D system, it turns out that only *time-delay* (TD) beam-forming, i.e. delaying the signals from the individual array elements by amounts equal to the relative differences in propagation time and then summing them, will make the aperture arrays function as intended. There are two main reasons for this:

- 1) The 3D system will routinely employ high-resolution, large bandwidth (up to 5 MHz) modulations. Reception of plasma lines at frequency offsets of up to  $\pm 15$  MHz ( $\pm 7\%$  of the centre frequency) will also be a routine feature. It is therefore absolutely necessary that the beam direction is independent of frequency, i.e. that the beam-former is *non-dispersive*. A true time delay system is indeed non-dispersive, but a system employing phasing introduced through a modulo- $2\pi$  phase-shifter is not; at large zenith angles, a  $\pm 7\%$  frequency change would shift the beam direction by several degrees!

- 2) The 3D arrays will be so large (more than 150 m diameter) that when very short pulses are used in order to achieve extreme range resolution, the backscattered signals from an individual pulse will not illuminate the whole array simultaneously. As an example, the 3D Core array has a diameter of 150 m while a 0.25 μs pulse has a spatial extent of 75 m; at a beam zenith angle of 45°, the pulse would at most cover only about 100 m of the array extent at once. If true time delay is applied to the individual signals before combining them, the pulse shape at the output of the beam former is correctly reconstructed. If modulo-2π phasing were used instead, the maximum time delay that could be applied would be equal to a period of the radar carrier frequency, about 4.3 nanoseconds! Severe pulse distortion will result and the received signals will be practically unusable.



**Figure 3:** Illustrating the principle of superposition.

Since the signal voltages from each array element are digitised individually at 80 – 90 Msamples/s, the time-delay function is conveniently implemented in the digital domain. For a given array, the maximum required delay is determined by the physical extent of the array and the maximum beam off-boresight angle. In the case of the 3D Core array, with a diameter  $\cong$  160 m and a maximum zenith angle of 40°, the maximum delay is about 370 ns, i.e. about 30 sampling intervals (assuming an 80 MHz sampling clock). At the same time, in order to achieve the beam pointing resolution of 0.06° specified in the 3D Performance Specification Document, it must

be possible to set the individual time-delays with a resolution better than 15 ps, i.e. about one-thousandth of a sampling interval! To handle this task “on the fly” at 80-90 MHz, we propose to use a system of *fractional-sample delay (FSD) filters*, i.e. all-pass FIR filters set up to work as interpolators. The FSD technique makes clever use of a consequence of the sampling theorem: *when a band-limited signal has been sampled at a rate exceeding twice the highest frequency present in the signal, the resulting sample series contains all information present in the original signal, so it is possible to reconstruct its instantaneous value at any past time by interpolation of a sufficiently long sequence of samples.* Interpolation of a sample series is commonly used in signal-processing systems and applications where a sample series must be re-sampled, either to a higher or to a lower sampling rate than the original one, but using interpolation to obtain a time-shifted replica of the original sample series at the original sampling rate seems to be much less common. A literature search shows that this technique has earlier been used in e.g. sonar [2,3], but as far as we have been able to determine, this is the first time that it is being proposed for use in a large research radar system.


A basic TTD beam-former is shown in Figure 4. It consists of two main parts, viz. a set of FSD units (blue) and a full-adder  $\Sigma$  (brown). The basic FSD function is realised as a generic FIR filter. One FSD unit is required per element antenna and beam. The filter coefficients  $c^m$ ,  $m = 1 \dots n$ , determine the filter group delay and must be determined independently for each element antenna and beam direction. By adding more taps to the filter, the integer part of the required delay can also be realised in the same structure.

An important bonus of this approach is that the FIR structure can also be used to simultaneously carry any number of other operations on the data streams, as long as these operations commute with the time-delay function. One obvious example is *aperture tapering*, which is frequently employed to suppress sidelobes at large off-boresight angles and may turn out to be required also in the 3D case; to accomplish the desired tapering, the coefficient sets for the individual array elements are simply multiplied by a set of scale factors representing the desired aperture distribution before being downloaded to their respective FSD units.

The data streams output from the FIR filters are all delayed by the correct amounts, such that summing them in the  $\Sigma$  block immediately generates the desired beam-formed data stream. *Multi-beaming can now be realised almost trivially by letting several beam-formers with different settings run in parallel on the same input data.* This feature will be particularly valuable at the receiver sites, where it will make it possible to observe a number of volumes along the transmitter beam truly simultaneously.

During the Design Study, several different approaches to synthesising FIR delay filters exhibiting almost perfect phase-linearity over a 30-MHz bandwidth have been tried. The LTU team has performed an extensive Matlab simulations to verify the validity of one approach [4]. 48-tap filters have been shown to provide the required fractional delay accuracy over a 30 MHz wide band while introducing very little amplitude ripple (less than a fraction of a dB).


The basic arithmetic operation required to realise an FIR filter is the Multiply-Accumulate (MAC). One MAC is required per filter tap and clock cycle, so a single 48-tap FSD filter running at 90 MHz requires  $4.32 \cdot 10^9$  MACs per second; the whole array (32000 individual receivers) requires close to  $1.4 \cdot 10^{13}$  MACs per second per beam.



**Figure 4:** Block diagram of a basic TTD beam-former, constructed from a number of FIR filters and a full-adder.

This massively parallel, computationally intensive task is an ideal application for *field-programmable gate array* (FPGA) technology. With the possible exception of application-specific integrated circuits (ASICs), no other technology can presently match the computational throughput / electrical power consumption ratio of FPGA. FPGAs can also be reconfigured in the field by reprogramming whenever improved algorithms are developed. This technology is therefore the logical choice for any application that is expected to remain in continual development over its entire lifetime, like the EI\_3D system.

To verify the TTD/FSD concept in practice, an FPGA-based beam-former has been developed for the Demonstrator array. A single Xilinx SX95T Virtex-5™ FPGA, implemented in 65-nm CMOS technology [5], was found to have sufficient raw parallel computing power to generate three simultaneous beams from the Demonstrator in a single chip. To speed up the design and implementation work, an evaluation board carrying the FPGA, a number of optical transceivers and most of the other required interface logic has been purchased. Figure 5 is a photograph of the evaluation board; a block diagram of the features available on the board is shown in Figure 6.

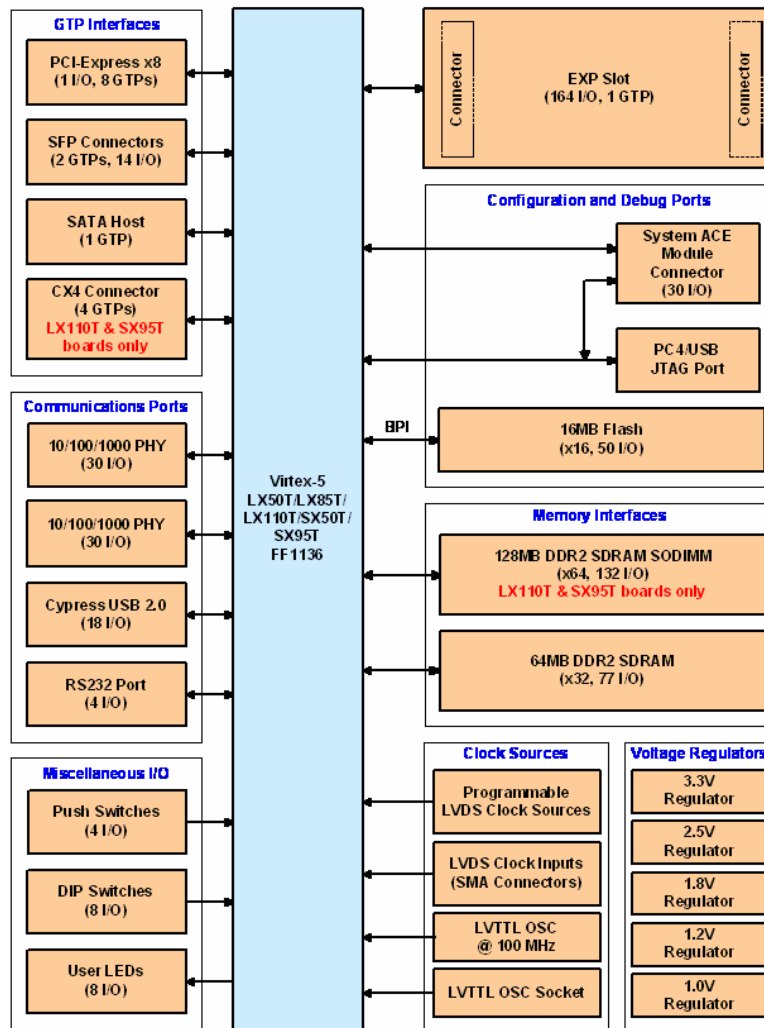


**Figure 5:** Photograph of the SX95T FPGA evaluation board used to implement the digital beam-former for the 3D Demonstrator array. Two optical transceivers can be seen on the left; an add-on board housing another four optical transceivers has been designed and manufactured at the EISCAT Sodankylä site.

The SX95T contains a total of 640 MAC blocks that can be clocked at up to 550 MHz. The multipliers are 25- by 18-bit  $2^s$  complement units and the accumulators can handle 48-bit precision. The chip also contains several application-specific hardware blocks (“sub-cores”) that can be used to unpack and de-serialise the stream of high-speed (1.6 Gb/s) data directly from the optical links to the computational core and to off-load the results onto other serial links. In an optimal configuration, the chip can theoretically deliver  $3.5 \cdot 10^{11}$  MACs per second, so something in the order of 50 SX95Ts could handle all computations required to form a single beam from the Core! In practice, a substantially larger number is likely to be required, since the real bottleneck now turns out to be the I/O capabilities rather than the MAC rate.

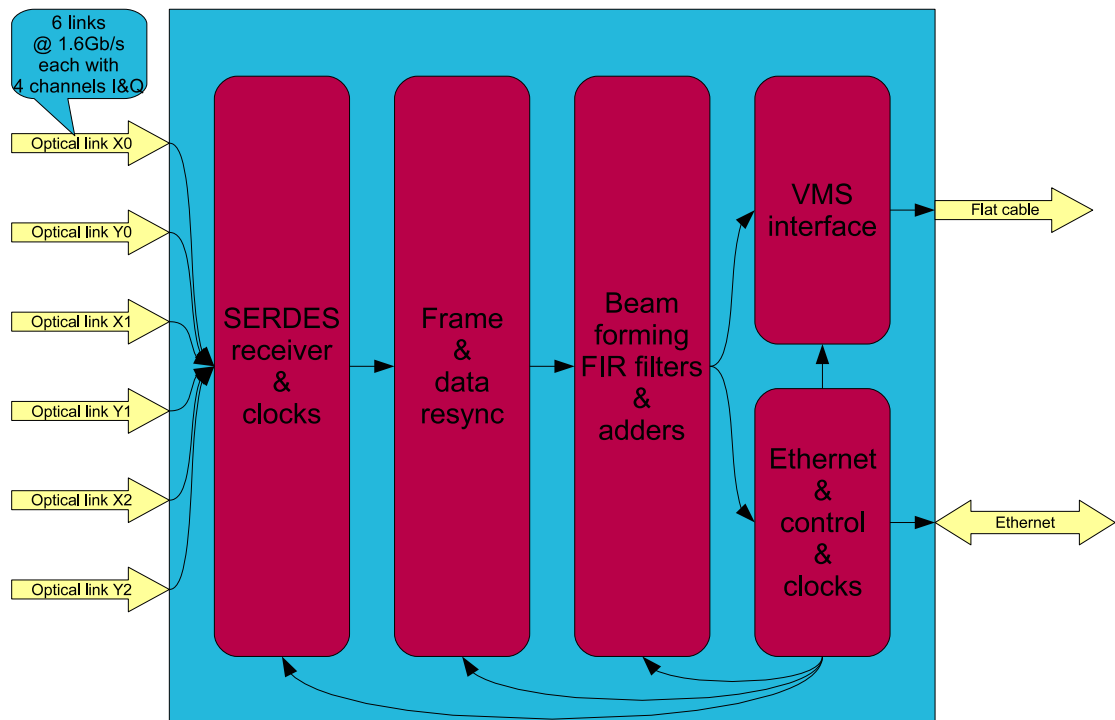
In the Demonstrator application, the FPGA computes three beams with X and Y polarization from an “apparent” array of twelve antennas (i.e. the 24 signals from the twelve array rows). Down-sampled I / Q signals from each antenna row are read-in at a data rate of 5 MHz. 48 FIR filters per beam are computed; with three beams, a total of 144 FIR 36-tap filters are thus running in parallel. To do this, the chip only needs to be clocked at 400 MHz.

To fully configure the three beams with correct delays, a total of 72 FIR coefficient filter sets have to be uploaded to the FPGA. Beam-forming sums are computed separately for the I and the Q data streams belonging to each beam, but since these use the same delay in their path from each antenna, identical FIR filter coefficients are used for the FIR pair. Loading, control and read-back are handled over a dedicated Ethernet connection to the host computer. Figure 7 shows a simplified diagram of the signal flow and resource usage on the SX95T evaluation board.



**Figure 6:** Block diagram of the resources available on the SX95T evaluation board

As shown in Figure 1, a convenient and practical way to arrange the beam-former system in the full-size arrays will be to build it as a hierarchical structure, where all data streams from the 49 X-Yagis connected to each equipment container are first fully beam-formed at the container level. This would probably require at least ten SX95T chips. The outputs from these “Level 1 beamformers” are directly useable as input to the next level, Level 2, which would essentially only consist of a set of full-adders, a small number of vector multipliers and the cross-correlators required for the coherency detection and triggering system. With this arrangement, fully beam-formed data streams from the individual containers are conveniently available to the coherency detector system and also immediately available for storage whenever the coherency threshold is exceeded.



**Figure 7:** Schematic diagram of the Demonstrator beam-former signal flow

### 2.3 Band-limiting and down-sampling; digital down-converters

The full bandwidth of the EI\_3D receiver system is specified at 30 MHz in order for the system to be able to simultaneously intercept signals at offsets of up to  $\pm 15$  MHz from the transmitter frequency. This frequency shift corresponds to plasma line returns from plasma with an electron density in the order of  $1.5 \cdot 10^{12} \text{ m}^{-2}$ . However, most of the information contained in the scatter signal is concentrated in the *ion line*, which is narrow (only a few kHz wide) and always located at most only a few kHz away from the transmitter centre frequency. The actual bandwidth of the received ion line signal is of course larger; it is equal to the extent of the convolution of the ion line spectrum with the spectrum of the transmitted probing signal and can therefore be as much as 5 MHz, if a high-resolution modulation is used. In a more typical case, it will be in the order of 200 kHz, or less than one percent of the full receiver bandwidth.

Many of the near-real-time data processing modules in the system, including the radar interferometry coherency detectors, the polarisation optimiser routines running at the receiving sites and the “standard quick-look” use only the ion line data. To reduce the computational burden to be carried by these applications, the signal processing chain includes a number of digital down-converter units (DDCs).

The functions performed by the DDCs are:

- 1) down-converting the band-segment of interest to baseband ,
- 2) band-limiting the resulting data stream as required by passing it through a programmable FIR filter, and
- 3) decimating the data stream from the FIR filter to a low rate, commensurate with the filter bandwidth.

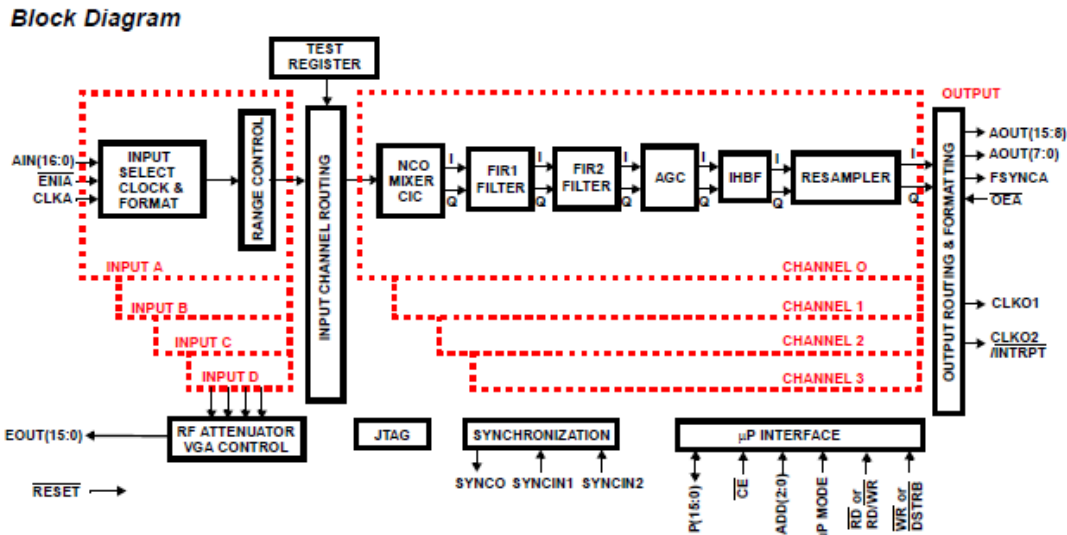
There will be two levels of DDCs. The data streams from the individual sub-arrays are processed in the first level; the outputs from these DDCs (DDC 1...DDC m) are used as the input signal for the coherency processors. The second DDC level (DDC A...DDC D) uses the fully beam-formed signal output from the Level 2 beam-former as its input and supplies band-limited data to the polarisation optimiser and the quick-look analysis routines discussed in section 4.

Fully integrated DDC subsystems are available from a number of manufacturers. One example is the ISL5416 four-channel ASIC used in the Demonstrator receiver system [6, 7, 8, 9]. Figure 8 shows a functional block diagram of this highly complex, extremely flexible component that can be operated at up to 95 M samples/s input rate and provides the following functions:

- Four Parallel 16-bit Fixed or 17-bit Floating Point Inputs
- Programmable RF Attenuator/VGA Control
- 32-Bit Programmable Carrier NCO with > 110dB SFDR
- 20-bit Internal Data Path
- Filter Functions
  - Multi-Stage Cascaded-Integrator-Comb (CIC) Filter
  - Two programmable FIR Filters (first up to 32-taps, second up to 64-taps)
  - Half Band Interpolation Filter
  - Re-sampling FIR Filter
- Overall decimation from 1 to >4096
- Digital AGC with up to 96dB of Gain Range
- Up to Four Independent 16-bit Parallel Outputs
- Serial Output Option
- 16-bit Parallel  $\mu$ P Interface
- 1.8V core, 3.3V I/O Operation

In the Demonstrator receiver, the ISL5416 chips are configured to band-limit the 80 Msamples/s data streams from the 24 digital front-ends to about  $\pm 0.5$  MHz and decimate them to 2.5 Msamples/s effective rate before transferring the data to the EISCAT UHF receiver.

Many FPGA realisations of DDCs also exist and can be had in the form of so-called “IP cores”, most of which are commercial products that are licensed to the end user for a fee. Either alternative would do the job; no recommendation as to which one to choose is made here as the market situation is likely to change substantially before the design must be finalised.



**Figure 8:** Functional block diagram of the ISL5416 digital down-converter

## 2.4 Coherency processors / triggers for interferometry

Radar interferometry and imaging will be standard features of the EISCAT\_3D radar. It will be possible to record the full-bandwidth raw data from a large number of previously selected sub-arrays ( $\approx 20$ ) in order to enable the subsequent processing of the data into 3-d volumetric images. But because the aggregate data rate to disk in this mode is some 10 - 20 times higher than when the system is running in its normal, beam-forming mode, sustained interferometry operation is not possible.

On the other hand, since running the system in interferometry mode is expected to be most meaningful when there are targets exhibiting a certain degree of coherency in the beam, an automatic *coherency processor* is included in the signal-processing system. Its job is to continually generate estimates of the degree of coherency between the signals received by pairs of separated sub-arrays at some time delay  $\tau$ . Whenever the coherency is found to meet or exceed some set threshold, the coherency processor issues an interrupt to the radar control system, requesting that the whole array should switch over into interferometry recording mode.

The coherency processor comprises three main parts, viz. the *cross-correlator/auto-correlator*, the *decision logic* and the *buffer memory*. The cross-correlator computes the product  $r_{jk}(n, t, \tau)$  of a pair of samples from the complex-valued, band-limited and down-sampled data streams from two sub-arrays  $j$  and  $k$  located the endpoints of an interferometry baseline:

$$r_{jk}(n, t, \tau) = s_j(n, t) \cdot s_k^*(n, t + \tau)$$

The products are averaged point-for-point over a number of consecutive radar cycles,  $n_{\max}$ , to generate an estimate of the cross-correlation or coherency between the two signals received by the two sub-arrays at a time delay  $\tau$ :

$$\langle r_{jk}(t, \tau) \rangle = \sum_{n=1}^{n_{\max}} s_j(n, t) \cdot s_k^*(n, t+\tau)$$

Before being input to the decision system, the  $\langle r_{jk}(t, \tau) \rangle$  estimates are normalised by dividing them by the square root of the product of the respective autocorrelations  $\langle r_{jj}(t, \tau) \rangle$  and  $\langle r_{kk}(t, \tau) \rangle$ , computed as above, in order to eliminate effects of variations in the absolute scattering cross sections.

The decision system may initially be realised as a set of hard logic comparators that generate an interrupt, instructing the array to switch over into full interferometry mode, whenever the coherency exceeds some preset threshold level. More elaborate algorithms, implemented in software, are also being considered. The correlators and comparators will probably be implemented in FPGA.

In order to make it possible to study the development leading up to the formation and detection of a coherent target, all data used to compute the  $\langle r_{jk}(t, \tau) \rangle$  estimates must be buffered for at least the number of radar cycles over which the averaging is performed, such that the data can be retrieved and post-processed whenever a coherency interrupt is generated. There are potentially at least two different ways to implement the buffering; either in a local RAM buffer or FIFO, or the data can be transferred into the central data store, where a dedicated ring buffer would have to be implemented on hard disk. The cleanest and least bandwidth-demanding alternative is probably to include a RAM buffer in the coherency processor; 10 seconds buffer time will require less than 2 Gbyte per data stream/baseline endpoint. Today, this is a trivial amount of RAM.

## 2.5 Quick-look analysis

The band-limited data from each fully formed beam in the EI\_3D system is fundamentally equivalent to the data output from the “channel boards” in the existing EISCAT VHF, UHF and ESR receiver back ends. The data processing software following the digital down-converters will therefore need to perform essentially the same tasks as the *lag\_wrap* and *decodump* processes running in the current EISCAT systems, that is, it should process the channel board data into complex autocorrelation coefficients, decode and re-shuffle these, and save the resulting autocorrelation *lag profiles* to disk.

The full lag profile set is one of the primary final data products from the system and will be securely archived. For control and monitoring purposes, selected parts of the lag profile data will be analysed in near-real time for physical parameter values (electron density, temperatures, Doppler velocity and polarisation) by e.g. the GUISDAP incoherent-scatter analysis software suite.

The *lag\_wrap* and *decodump* routines are easily integrated into any UNIX/Linux environment. They are initialised from simple scripts, describing how the raw data are arranged and how the data manipulation primitives should be applied. They are written in multi-threaded fashion, using POSIX pthread, and can thus immediately make full use of the computing power of bigger shared memory processor (SMP) machines if and when required; their performance will scale in proportion. Since these routines and the other current EISCAT data processing software already meet the 3D post-DDC processing requirements and can be readily re-used, we conclude that no major development of new data processing software is needed to implement a quick-look analysis function into the 3D system.

## 2.6 Polarisation matching

Although not included in the data flow diagram in Figure 1, the real-time merging of orthogonal polarisation components offers a possibility to reduce the data rate from the arrays by a factor of two and should therefore be briefly mentioned here:

To recover all energy and information present in the scattered signal from a particular direction on the sky, two co-linear beams must be formed, one for each of the two orthogonal polarisations. When the polarisation of the received wave is known, the two resulting 80-MHz data streams can be combined into a single data stream containing all the signal energy without losing any information about the target state [10]. The merging operation, which immediately reduces the data rate to disk by a factor of two, can be performed by a vector rotation and subsequent addition, two operations requiring only a handful of MAC blocks and therefore representing a totally negligible computational load in comparison to the beam-forming task; the merging technique is therefore very appealing from a storage and bandwidth conservation point of view.

At the Core array, it will be possible to arrange matters such that the received polarisation is always close to circular, making beam merging a straightforward process. The biggest gains can however be reaped at the remote sites, where five or more simultaneous beams will be in operation. Here, the received polarisations are time-varying functions of the amount of Faraday rotation suffered by the scattered signals on their way from the scattering region to the receiver, and therefore not known *a priori*; the actual polarisations must be determined from the respective received signals and the corresponding scaling factors must be computed and fed back into the FPGA system. These operations require a fair amount of statistical processing of the incoming data in near-real time, using the so-called *subspace tracking* technique. For this purpose, a high performance, general-purpose computing platform must be attached to the signal processor, taking its input from a second-level DDC.

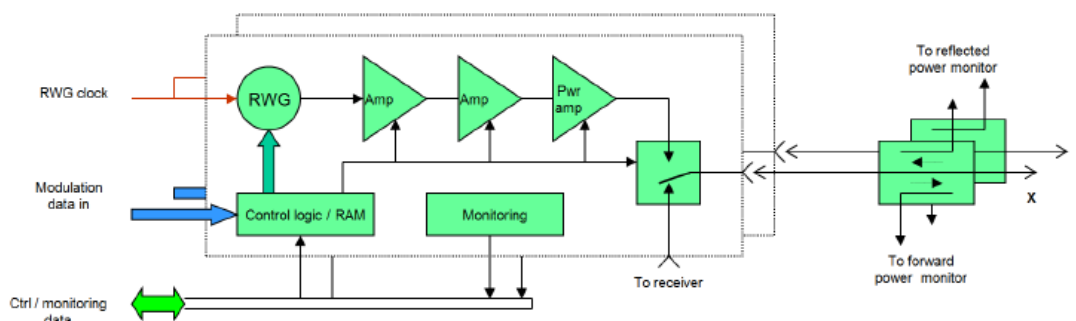
### 3. Transmitter system signal processing

#### 3.1 The basic transmitter module

The basic 3D transmitter performance requirements, as listed in the “EISCAT\_3D Performance Specification Document”, are:

- Centre frequency: between 220 – 250 MHz, subject to allocation
- Peak output power:  $\geq 2$  MW
- Instantaneous  $-1$  dB power bandwidth:  $\geq 5$  MHz
- Pulse length: 0.5–2000  $\mu$ s
- Pulse repetition frequency: 0–3000 Hz
- Modulation: Arbitrary waveforms, limited only by power bandwidth

In an ideal EI\_3D transmitter system, these requirements should be met by fitting a self-contained digital arbitrary waveform/linear RF power amplifier unit (Figure 9) to each array element. The resulting system configuration would be capable of handling any arbitrary combination of phase and amplitude modulation, while at the same time providing the hardware functionality required for implementing both the time-delay beam-steering function and, if required, array aperture tapering on transmit. The arbitrary-waveform capability could also be used to suppress out-of-band emissions by applying tailored spectrum masks to all transmissions, something which might turn out to be required at all times in order to prevent interference to the active DAB channels immediately below and above the 3D spectrum slot. This “ideal” transmitter system can in fact be implemented now, using components available off-the-shelf.



**Figure 9:** Schematic diagram of an element transmitter, comprising two identical units, one for each of the two polarisations. Each unit contains a digital random (and/or arbitrary) waveform generator/upconverter (RWG), two intermediate power amplifier stages, a 300+ watt power amplifier, a transmit-receive switch, a RAM bank and control and monitoring logic. Depending on the degree of mutual coupling between neighbouring element radiators, circulators and waster loads may have to be fitted between the power amplifier and the transmit/receive switch.

### 3.2 RF waveform generation requirements

According to presently available information, the EISCAT\_3D transmitters will be operating in the 230 – 236.7 MHz frequency band on a non-interference basis; the entire band will be available for transmissions and no restrictions will be put on the in-band waveforms. Out-of-band radiation will however have to be strongly suppressed in order to meet the non-interference requirement; this may require the application of spectral masks. The EI\_3D arbitrary-waveform generator must meet all these requirements.

Arbitrary-waveform signals at high RF frequencies are commonly generated by *up-conversion* of a baseband signal. An I/Q carrier signal at the desired RF frequency is fed into an I/Q mixer, where it is multiplied with the desired arbitrary I/Q baseband signal. The mixer output is a replica of the baseband signal, translated in frequency by an amount equal to the magnitude of the carrier frequency.

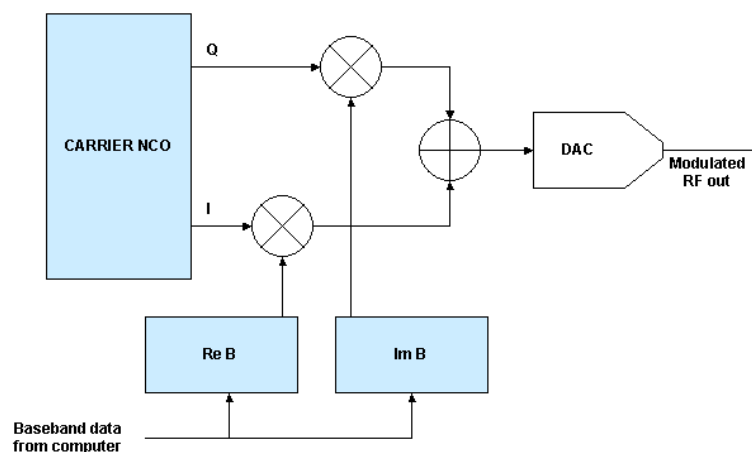
The up-conversion technique works equally well both in the time-continuous analogue domain and in the discrete-time digital domain, provided phase quadrature and amplitude balance are maintained throughout. In the digital realisation (Fig. 10), the carrier signal and the modulating signal are both complex-valued sample streams:

Carrier signal  $C_n = k_C \exp(i\omega\tau), n = 1 \dots$

Baseband (modulating) signal  $B_n = B(n), n = 1 \dots$

and the mixer is realised as a dual multiplier – full adder, generating the sequence  $S_n$  :

$$S_n = k_C [\cos(i\omega\tau) \cdot \text{Re } B(n) + \sin(i\omega\tau) \cdot \text{Im } B(n) ]$$



**Figure 10:** Functional block diagram of a digital arbitrary-waveform generator

The B data stream must be clocked into the multipliers at least at the baseband Nyquist rate, but can of course be clocked in at several times this rate to enjoy the usual benefits of oversampling, if memory bus bandwidths and control computer speed allow. After digital-to-analogue conversion, the mixer output signal is an analogue RF voltage signal that can be used to drive a transmitter power amplifier system.

In the EI\_3D case, the maximum permissible modulation bandwidth is 6.7 MHz, limited by the frequency allocation. An up-conversion exciter system can therefore be designed with comparatively moderate B memory speed and baseband data bus bandwidth; feeding data into the multipliers at 30 MHz rate will already permit four-fold oversampling.

Since both the amplitude and the phase of the RF signal can be arbitrarily modulated, no special RF switch is required in order to turn the RF drive signal on at the beginning of a pulse and off at the end; when no RF is to be transmitted, the B data is simply set to all zeros. Assuming a 30 MHz baseband data rate, the resulting pulse-timing resolution is about 33 ns.

Array aperture tapering while in transmitting mode will be easily implemented by scaling the B data fed to the individual exciters through multiplication by an amplitude distribution function before downloading. In its simplest form, this function is scalar, but a complex-valued function will be accommodated just as easily; this will allow e.g. the transmission of signals with orbital angular momentum  $L \neq 0$ .

The waveforms to be transmitted must of course be designed, computer-generated and downloaded to the B memory banks before they can be executed on the radar. The design and generation can be conveniently done with any one of a number of standard software tools, e.g. the Matlab Signal Processing Toolbox.

### **3.3 Beam steering and beam-forming**

The same arguments as those raised in Section 2.2 apply also to the transmitting case; only *time-delay* (TD) beamforming will work.

In practice, this means that a FSD system, functionally equivalent to that implemented in hardware on the receiver side, must be used to interpolate the modulating waveform (the B data stream) to the required time resolution before it is uploaded to the individual exciter units.

The main difference between transmit and receive modes is that in the transmit mode, the modulating waveforms are designed rather than received and so known in advance. It can also be assumed that in most cases a common modulating waveform is going to be used for all elements in the array, such that the interpolation can be performed ahead of runtime, e.g. by software linked to the waveform design programme package.

Linearity across the entire modulation bandwidth will again be required in order that all spectral components are radiated into the same direction, so similar considerations will apply to the design of the FIR interpolating filter. However, if the exciter system can be designed to allow fourfold oversampling of the baseband signal, it should be possible to relax the filter-design boundary conditions substantially.

### **3.4 An example of a commercial DUC – the AD9957**

The arbitrary-waveform generator building block is clearly a key component. Since tens of thousands of these will be required, they must be available at low cost for the scheme to be practically and economically feasible. While the required functionality can be provided by a laboratory vector signal generator/random waveform generator attachment, such as is available from e.g. Rohde & Schwarz, this comes with a price tag in the order of 50 kEUR and is therefore out of the question; another solution must be found.

However, a limited offering of affordable digital up-converter mixed-signal integrated circuits incorporating the desired functionality was in fact available already at the beginning of the Design Study in 2005. At that time, no device was yet fast enough to be directly useable as the exciter for a 3D transmitter, but steady progress toward higher clock frequencies was being made. Since it was expected that devices capable of directly generating a modulated carrier at 230 MHz would be available by the end of 2008, the project team has kept the development in this area under observation. A silicon subsystem offering all the required performance is now indeed available: the recently released Analog Devices, Inc. AD9957. This device provides the following features:

- 1 GSPS internal clock speed (up to 400 MHz analog output)
- Integrated 1 GSPS 14-bit DAC
- 250 MHz I/Q data throughput rate
- Phase noise  $\leq -125$  dBc/Hz (400 MHz carrier @ 1 kHz offset)
- Excellent dynamic performance >80 dB narrow-band SFDR
- 8 programmable profiles for shift keying
- SIN(x)/(x) correction (inverse sinc filter)
- Reference clock multiplier
- Internal oscillator for a single crystal operation
- Software and hardware controlled power-down
- Integrated RAM
- Phase modulation capability
- Multichip synchronization
- Interpolation factors from 4× to 252×
- Interpolation DAC mode
- Gain control DAC
- Internal divider allows references up to 2 GHz
- 0 dBm output into 50 ohm without exceeding voltage compliance rating

A functional block diagram of the AD9957 is given in Figure 11. Full specifications are available in [11].

Since the AD9957 can generate the EI\_3D carrier frequency (230 – 236 MHz) directly, accept modulation I/Q data at up to 250 MHz and deliver an output power of 0 dBm into 50  $\Omega$ , it can indeed do everything required of the arbitrary-waveform generator part of Figure 9. A practical and affordable solution to the transmitter exciter design problem thus appears to be within reach. Figure 13 gives an example of the kind of performance that can be expected at  $\approx$  220 MHz when the device is clocked at 1 GHz. In the single-tone case shown in the two upper panels, the carrier is suppressed by  $-80$  dB relative to the desired sideband; the unwanted sideband is suppressed by more than  $-100$  dB and invisible in the noise. In the lower two panels, a wideband (7.8125 MHz) baseband signal with a moderate amount of spectral shaping modulates the carrier; at  $\pm$  5 MHz, the spectrum has dropped off by more than 65 dB relative to the average in-band level. It is clearly possible to use this device to generate anything from a monochromatic carrier to a pseudo-random signal that uniformly fills the entire bandwidth allocated to the EI\_3D system with spectral power.

AD9957 is now available off-the-shelf and it is also relatively inexpensive; current pricing is around USD 25 in 1000+ quantities! Unfortunately, because the device became commercially available at such a late stage of the 3D project, there has so far been no possibility to evaluate its capabilities in full. However, two AD9957 evaluation boards (Figure 12) have been ordered and delivered. Further work will concentrate on mating the AD9957 to the clock distribution and data transmission strategies developed in WPs 4 and 12 and to incorporate it in the test bed (see below) to enable full-bandwidth, random-waveform testing of the power amplifier chain.






Figure 6. 15.625 kHz Quadrature Tone, Carrier = 222 MHz, CCI = 16,  $f_s = 1$  GHz




Figure 9. Narrow-Band View of Figure 6 (with Carrier And Lower Sideband Suppression)

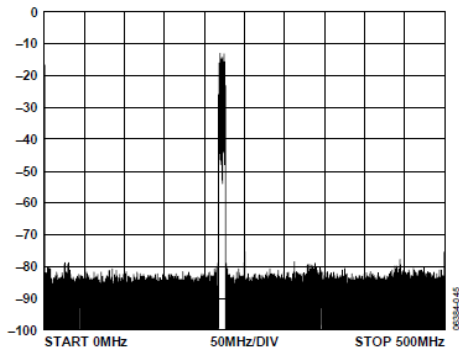


Figure 12. QPSK, 7.8125 Msymbols/s, 4x Oversampled Raised Cosine,  $a = 0.25$ , CCI = 8, Carrier = 222 MHz,  $f_s = 1$  GHz




Figure 15. Narrow-Band View of Figure 12

**Figure 13:** Examples of signal spectra generated by the AD9957 (courtesy Analog Devices). The upper two panels show viz. a wideband and a high-resolution power spectrum of a 15.625 kHz tone, quadrature modulated onto a 222.000 MHz carrier. In the lower two panels, a 7.8125 Msymbols/s spectrally shaped baseband signal is modulated onto the same 222.000 MHz carrier, producing a  $\approx 8$  MHz wide power spectrum. At  $\pm 5$  MHz, the spectrum has dropped off by more than 65 dB, illustrating the spectrum-shaping effect of the raised-cosine window applied to the baseband signal.

## References

1. [http://e7.eiscat.se/groups/EISCAT\\_3D\\_info/d42](http://e7.eiscat.se/groups/EISCAT_3D_info/d42)
2. Murphy, N.P., A. Krukowski and I. Kale, Implementation of a Wide-Band Integer and Fractional Delay Element, *Electronics Letters*, 30, No. 20, pp. 1658-1659 (1994)
3. Murphy, N.P., A. Krukowski and A. Tarczynski, An Efficient Fractional Sample Delayer for Digital Beam Steering, *Proceedings of ICASSP '97* (1997)
4. G. Stenberg, J. Borg, J. Johansson, and G. Wannberg, *Simulation of Post-ADC Digital Beam-Forming for Large Area Radar Receiver Arrays*, in *International RF and Microwave Conference*, 2006.
5. [http://www.xilinx.com/support/documentation/data\\_sheets/ds100.pdf](http://www.xilinx.com/support/documentation/data_sheets/ds100.pdf)
6. <http://www.intersil.com/data/rd/commlink/isl5416/isl5416info.pdf>
7. <http://www.intersil.com/cda/deviceinfo/0,1477,ISL5416,00.html#longdesc>
8. <http://www.intersil.com/cda/deviceinfo/0,1477,ISL5416,00.html#data>
9. <http://www.intersil.com/data/fn/fn6006.pdf>
10. [http://e7.eiscat.se/groups/EISCAT\\_3D\\_info/D9.3](http://e7.eiscat.se/groups/EISCAT_3D_info/D9.3)
11. [http://www.analog.com/static/imported-files/data\\_sheets/AD9957.pdf](http://www.analog.com/static/imported-files/data_sheets/AD9957.pdf)