

ADC Sample Rate and Preliminary Design for a Full-RF ADC Post-Processor

Ver. 0.1

Steve Ellingson*

September 11, 2007

Contents

1	Introduction	3
2	Design Concept	4
3	Preliminary Implementation	9
4	Issues to Address in Future Versions of this Document	13
5	Document History	14

*Bradley Dept. of Electrical & Computer Engineering, 302 Whittemore Hall, Virginia Polytechnic Institute & State University, Blacksburg VA 24061 USA. E-mail: ellingson@vt.edu

List of Figures

1	Signal flow.	6
2	Frequency plan and processing through $F_S/4$ -shift-left.	7
3	Frequency plan and processing (continued from Figure 2).	8
4	Low pass FIR filter characteristics (screenshot from the design software). Not shown on this screen is the number of taps, which is 40. Also not shown here is the fact that a Blackman window has been applied to the coefficients, which accounts for the very low sidelobe level. The horizontal axis of the frequency response plot is in units of the input sample rate, so 0.3 corresponds to 59 MHz at 196 MSPS.	11
5	Summary of the design/synthesis flow summary (screenshot from the design software).	12

List of Tables

1	Characteristics of the FPGA selected for the preliminary implementation.	10
---	--	----

1 Introduction

This document presents a preliminary design for the first stage of digital processing in the LWA signal chain. Context for this work can be found in the System Architecture document [1], currently in Version 0.5. This document addresses analog-to-digital converter (ADC) sample rate, ADC location, and details of the first stages of digital signal processing. This document resolves the “open question” from [1] about the location of the ADC: It will be located in the DP1 subsystem as opposed to the ARX. A full-RF digital “postprocessor” for the ADC is described, consisting of a delay line, quadrature downconverter, additional filtering to define the desired 10–88 MHz (with respect to sky frequency) passband, and decimation to reduce sample rate commensurate with the new bandwidth. It is convenient to perform these functions in a postprocessor immediately following the ADC, as the same processing would otherwise be required to be repeated in each of the BFUs, and an FPGA is required for data transmission format conversion anyway. The postprocessor design is shown to be implementable in an FPGA costing less than US\$50 (quantity=1).

2 Design Concept

Figure 1 shows the signal flow from ADC through the output of the first FPGA. It is envisioned that both the ADC and FPGA will be located in the DP1 subsystem. In this scheme, the ADC digitizes the output from an ARX (located in the ASP subsystem) at a sample rate $F_S = 196$ MSPS. This sample rate turns out also to be quite convenient for subsequent multirate processing, as will be demonstrated below. The ADC selection process [2] yielded several suitable candidates, with either 10 or 12 bits. The basis for selecting $F_S = 196$ MSPS is illustrated in Figure 2(a), which shows that digitization at this rate causes the 88–108 MHz FM broadcast band to alias onto itself, with no ingress whatsoever into the 10–88 MHz tuning range. This loosens the requirement for anti-alias filtering, since the lowest frequency that can then alias into the “desired” tuning range is 108 MHz, and the lowest frequency that can alias into the “required” tuning range is 116 MHz. Suitable filters will not be difficult to design, especially since the spectrum above 108 MHz appears to much less of a problem than the FM band in most locations.

The output of the ADC is assumed to be either 10-bit or 12-bit parallel LVDS (consistent with [2]) at 196 MSPS. We shall assume the worst case (from the perspective of digital design) of 12 bits. Since there are 512 ADCs for a 256-stand station, interconnects deserve careful attention. If the ADCs are located in ARXs and separated from the DP1 subsystem, then this poses a serious problem. Implementing this as 512 12-bit parallel data connections is not out of the question technically, but it will pose quite a challenge in terms of space and organization within the electronics shelter, and poses potential signal integrity, electromagnetic interference (EMI), and maintenance headaches. This could be reduced to 512 serial outputs; in this case the speed is 2.352 Gb/s per connection. This is no problem technically; in fact, the nodal interconnections in the reconfigurable computing cluster of the Eight-meter-wavelength Transient Array (ETA) [3] operate at about this rate over 3 m cables. This scheme is far more reasonable from both a signal integrity and EMI perspective, as long as a suitable electromechanical interface standard such as Infiniband (the scheme used in ETA) is employed. However, the maintenance issue remains and the cost per connection is potentially quite high – at least tens of dollars per link for an Infiniband solution, for example. The alternative is to place the ADC in the DP1 subsystem, preferably on the same circuit board as the first FPGA in the signal flow. In this case, the signal integrity and maintenance problems become relative easy. The EMI generated is probably less, however the risk of contamination of data by EMI is probably greater due to the ADC’s proximity to digital hardware. However, there is good anecdotal evidence that this risk is justified. For example, ETA uses ADCs implemented on the same board as a large FPGA; also the Berkeley system currently favored by the ATA and PAPER projects use ADCs on or directly connected to FPGA boards. If we follow this approach for LWA, the interconnect between ASP and DP1 becomes a single coaxial cable (512 cables total for a station). These are inexpensive (less than \$10 per connection) and relatively easy to organize.

Thus, the ADC will be located in the DP1 subsystem, on the same circuit board as the next major component in the signal flow, which will be the FPGA postprocessor. Collectively the ADC and FPGA-postprocessor will be referred to as “the digitizer” and the set of all digitizers will appear as a new sub-subsystem of the DP1 subsystem in the next version of the System Architecture Document. The postprocessor will implement multiple functions, as shown in Figure 1. Some of these will be diagnostic features that will be transparent to the normal operation of the system but will necessary or desirable for development¹. The first stage of consequence to normal system operation is a length-64 FIFO sample buffer implementing a system-configurable time delay, called out in [1]. In this case, 196 MSPS corresponds to 1.53 meters per sample, so a length-64 first-in first-out (FIFO) can implement the longest possible differential delay across the 100 m station aperture. Furthermore, this is just a fraction of a wavelength at the highest required frequency of operation (80 MHz, corresponding to 3.75 m wavelength), making it feasible to use this mechanism as the delay portion of a crude delay-and-sum beamformer

¹No attempt will be made to address these in this version of this document.

should this ever be desirable.

In the same FPGA, it is useful to implement any processing which is common to any sub-subsystems appearing in the signal flow; BFUs and DRXs in particular. Here, it is proposed to reduce the data to complex baseband (i.e., “I-Q”) form and to reduce the bandwidth and sample rate accordingly. The method shown here uses an “ $F_S/4$ -shift-left”, followed by a low pass filter, followed by decimation by a factor of 2. It should be emphasized at this point that this particular sequence of operations is shown only for the purposes of clearly indicating intent, and is surely NOT the most efficient way to implement the desired processing. It is quite likely that combinations of these operations could be done more efficiently using multirate techniques with considerable reduction in required hardware resources; see e.g. [5]. Proceeding with this understanding, we now describe the processing occurring at each step.

$F_S/4$ -shift-left: This is a form of tuning which exploits the fact that a spectral shift equal to one-fourth of the sample rate (49 MHz in this case) is equivalent to multiplication by $+1$, $+j$, -1 , $-j$, $+1$, ... [4]. This reduces the operation of quadrature downconversion to a process of demultiplexing the real-valued samples into I and Q streams, and alternating signs as necessary. No multiplications are required, which makes this extremely attractive especially when it is required to minimize FPGA resource requirements. The spectrum before and after the $F_S/4$ -shift-left is shown in Figures 2(b) and (c), respectively. Note that the 10–88 MHz tuning range now occupies the range –39 to +39 MHz.

Low-pass filter: In order to reduce the sample rate, it is first necessary to suppress portions of the spectrum that will alias into the desired spectrum after decimation. In this design we are fortunate that the folding frequency (49 MSPS) after decimation again falls exactly in the center of the (shifted) FM broadcast band; furthermore, the same is true of the (now shifted and double-sided) DC–10 MHz band. Thus, the lowest frequencies which need to be considered from an anti-aliasing perspective is 59 MHz (which aliases to the top edge of the shifted “desired” tuning range) and 67 MHz (which aliases to the top edge of the shifted “required” tuning range). The spectrum before and after low-pass filtering is shown in Figures 2(d) and 3(a), respectively.

Decimation by two: The spectrum after decimation by 2 is shown in Figure 3(b). The sample rate is now 96 MSPS, complex.

It should be noted that some degree of RFI mitigation could and perhaps should be integrated into the above signal flow. If effective, RFI mitigation could dramatically reduce the number of bits/sample from greater than 10 to perhaps as few as 4. This would have enormous benefits in terms of simplifying design.

The output of the FPGA postprocessor should be packaged in a form convenient for transport to the next sub-subsystem. According the current system architecture, this will be a BFU and the method of transport is a key feature of the DP1 daisy chain. Assuming the BFU exists on the same circuit board or nearby (e.g., on a daughterboard), there are several options. First, we might consider moving the output in parallel form. The efficacy of this depends on the number of bits per sample. If this number is greater than 4–8, then we run the risk of needing more expensive FPGAs (in order to obtain the desired number of I/O pins), and also face some challenges in circuit board trace routing. A better approach in this case is probably to exploit the high data rate available using the native LVDS capability now provided in most FPGAs. LVDS links can run reliably at speeds up to 800 Mb/s. Thus, one way to minimize the number of required inter-FPGA links is the technique shown in Figure 1. Here, the complex data stream is assumed to be 24 bits wide (i.e., 12-bits “I” + 12-bits “Q”). This is reduced to a 12-bit-wide data stream at 196 MHz simply by multiplexing the “I” and “Q” streams. Although many suitable FPGAs have sufficient numbers of LVDS channels, circuit board design is greatly simplified by reformatting this 12-channel output as 3 channels, each 784 Mb/s – a nice fit to the typical upper bound for LVDS of about

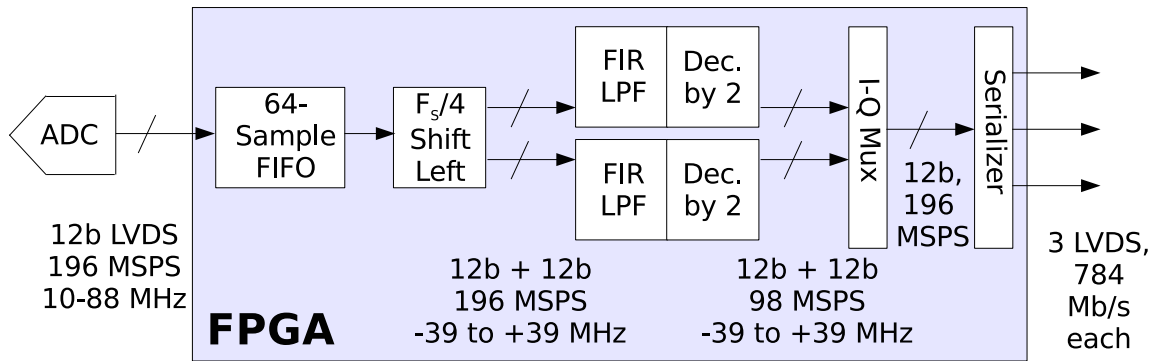




Figure 1: Signal flow.

800 Mb/s. Thus, it is suggested that the data be moved off the FPGA post-processor using 3 serial LVDS links. This then becomes a feature of the DP1 daisy chain specification. Taking into account the 12 inputs required for interface with the ADC, plus additional LVDS channels for clock input and output, then we will need an FPGA with 17 LVDS serial ports. Such FPGAs are commonly available at reasonable cost, as will be demonstrated in the next section.


A case study which is useful as a tutorial in interface design of this type, and which may be useful as an example of what is possible, is available in [6].



(a) Spectrum at input to ADC. *Dark green*: 20-80 MHz “required” tuning range. *Light green*: 10-88 MHz “desired” tuning range. *Red*: 88-108 MHz FM broadcast band. *Orange*: DC-10 MHz; potential source of strong signals (esp. shortwave and AM broadcast). *Yellow*: The 108-196 MHz alias band – signals in this range alias into the DC-88 MHz band after digitization at 196 MSPS.



(b) Spectrum at output of ADC.



(c) Spectrum at output of $F_S/4$ -shift-left. Yellow alias bands now omitted for clarity.

Figure 2: Frequency plan and processing through $F_S/4$ -shift-left.





Figure 3: Frequency plan and processing (continued from Figure 2).

3 Preliminary Implementation

To confirm the efficacy of the scheme proposed in the previous sections, a candidate FPGA was selected and firmware was developed for it. The selected FPGA is the Altera EP3C25, which is one of the “Cyclone III” family of devices [7]. The characteristics of the specific part selected are summarized in Table 1.

Firmware was developed in Altera’s Quartus II software, Ver. 7.1. 12-bit ADC input and an input clock are accepted through LVDS pins and processed through an $F_S/4$ -shift-left. The output feeds a pair of low-pass FIR filters developed using Altera’s FIR Compiler IP core, Version 7.1. The parameters of the FIR design are summarized in the screenshot shown in Figure 4. Note that rejection greater than 70 dB is achieved for frequencies greater than 59 MHz. Also note that this design includes an integrated decimate-by-2, resulting in a considerable performance improvement with respect to the concept design described earlier where these are separate operations. 24 bits (i.e., 12 bits “I” + 12 bits “Q”) are selected from the FIR outputs and muxed onto 12 LVDS output channels in order to force synthesis of the relevant sections of the filter and to ensure a conservative assessment of timing efficacy.

A summary of the completed design flow is shown in Figure 5. Note that the design consumes less than 30% of the logic elements on the part, allowing considerable flexibility in adding additional features. This design achieves timing closure for input clock rates up to 219 MHz; thus the desired 196 MHz is a comfortable fit.

It should be noted that an input FIFO delay line was not implemented in this design although ample memory and logic elements certainly exist for that feature. Also, the serializers and 3-channel-output implementation called out in the previous section was not implemented primarily because it results in a less conservative estimate of resources and timing feasibility; however these too are easily implementable with the remaining resources.

Vendor	Altera
Part Number	EP3C25F324C8 Short form: "EP3C25"
Device Family	Cyclone III
Logic Elements	24,624
Memory	594 Kb
Multipliers	66, each 18 × 18 bit
PLLs	4
User I/O Pins	216
Differential Channels	71
Package	324-pin FBGA
Price	US\$49.30, quantity=1, via Altera on-line store (www.altera.com)
Availability	175 available immediately at above source (No other sources investigated)

Table 1: Characteristics of the FPGA selected for the preliminary implementation.

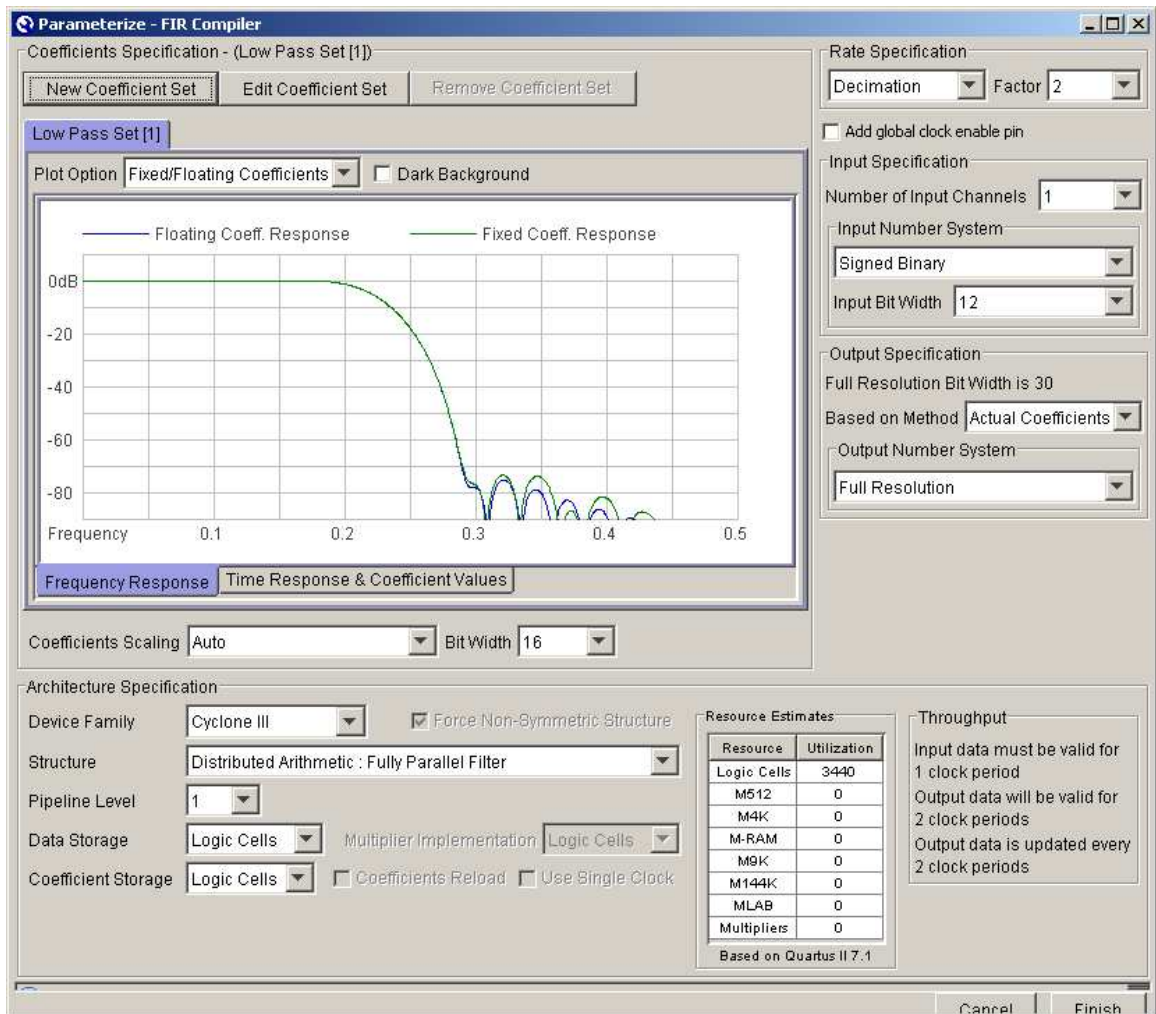


Figure 4: Low pass FIR filter characteristics (screenshot from the design software). Not shown on this screen is the number of taps, which is 40. Also not shown here is the fact that a Blackman window has been applied to the coefficients, which accounts for the very low sidelobe level. The horizontal axis of the frequency response plot is in units of the input sample rate, so 0.3 corresponds to 59 MHz at 196 MSPS.

Flow Summary	
Flow Status	Successful - Mon Sep 10 21:12:29 2007
Quartus II Version	7.1 Build 178 06/25/2007 SP 1 SJ Full Version
Revision Name	DR1
Top-level Entity Name	DR1
Family	Cyclone III
Device	EP3C25F324C8
Timing Models	Preliminary
Met timing requirements	Yes
Total logic elements	7,176 / 24,624 (29 %)
Total combinational functions	6,459 / 24,624 (26 %)
Dedicated logic registers	5,170 / 24,624 (21 %)
Total registers	5170
Total pins	66 / 216 (31 %)
Total virtual pins	0
Total memory bits	0 / 608,256 (0 %)
Embedded Multiplier 9-bit elements	0 / 132 (0 %)
Total PLLs	1 / 4 (25 %)

Figure 5: Summary of the design/synthesis flow summary (screenshot from the design software).

4 Issues to Address in Future Versions of this Document

1. Consider more aggressive digital filtering of the FM band before decimation. Even though it is not an aliasing threat, reducing residual levels allows the number of bits per sample to be reduced.
2. Consider an ADC clip detector/counter at input to FPGA.
3. Use whatever RAM is available on the FPGA to make a short capture buffer (perhaps through JTAG) to facilitate development. Sufficient memory exists for capture of up to tens of thousands of samples. Could be extension of the FIFO delay line buffer.
4. If the ADC uses an SPI control port, the FPGA described here would be the logical place to connect it to.
5. More explicit description of $F_S/4$ processing; show the math. Alternatively, work out a closer-to-optimum multirate implementation.
6. Work out optimum length of low-pass filter and number of bits for data and coefficients.
7. Repeat Cyclone III synthesis including FIFO filter, actual output mux, etc.

5 Document History

- This is Version 0.1, which is the first version released.

References

- [1] S. Ellingson, "LWA Station Architecture Ver 0.5," August 28, 2007.
- [2] S. Ellingson, "ADC Selection," LWA Memo 98, August 31, 2007.
<http://www.phys.unm.edu/~lwa/memos/>.
- [3] "Eight-meter-wavelength Transient Array (ETA)," project web site,
<http://www.ece.vt.edu/swe/eta>.
- [4] J. Tsui, *Digital Techniques for Wideband Receivers*, Artech House, 1995. See Ch. 8.
- [5] R.E. Crochiere and L.R. Rabiner, *Multirate Digital Signal Processing*, Prentice-Hall, 1983.
- [6] Altera Corp., "Stratix Devices & Fujitsu MB86064 DACs," Application Note AN316, July 2003.
<http://www.altera.com>.
- [7] Altera, Inc., *Cyclone III Device Handbook*, July 2007. This is a two volume set: Currently Vol. 1 is in Version 1.1, Vol. 2 is in Version 1.3.
<http://www.altera.com>.