Transient Buffer – Wideband (TBW) Preliminary Design Ver. 0.1

Steve Ellingson*

November 11, 2007

Contents

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

1 Summary

This document presents a preliminary design for the LWA "transient buffer – wideband"(TBW) level-2 subsystem of the DP1 level-1 subsystem. The purpose of the TBW is to capture the complete output from all antennas for as long as possible, and then to deliver the captured data to the DAC for further processing or transport. Context for this work can be found in the System Architecture document [1], currently in Version 0.6, which defines the TBW and the DP1. The TBW accepts the "full RF" (78 MHz bandwidth sampled at 98 MSPS) output from 512 DIGs upstream in the DP1 daisy chain. Each of these 512 inputs represents one antenna; that is, one polarization from one stand in the station's antenna array. The DIG subsystem is described in [2] and the associated "DP1 daisy chain" 4-lane LVDS interface is described in [3].

Two TBW designs are described here. The first is derived through small firmware-only modifications to the BFU hardware, which is currently described in [4]. It is assumed that readers are familiar with that document. The cost of this minimal (41 μs) TBW, capable of capturing $N_{bs} = 28$ bits/sample per antenna, will be similar to that for a single BFU; i.e., \$20,000–\$33,000 per [4]. The second TBW design requires modification of the BFU-based hardware design to incorporate 4 PC133-based SDRAM chips per FPGA. This scheme permits captures of up to 85.6 ms, but with N_{bs} reduced to 16. This scheme adds about \$5,000 to the per-TBW cost and will probably require modification of BFU-based design in such a way that the TBW must always be the last subsystem in the DP1 daisy chain. In either scheme, capture time can be extended by reducing N_{bs}, with the new capture time increasing linearly with the reduction in N_{bs}.

It should be noted that neither of the designs described here will be satisfactory for most transient science applications, due to the small capture time. However, the proposed designs are suitable for use in system development/integration, and also in support of calibration of the station's antenna array and for per-antenna diagnostic operations. A limited but significant all-sky transient monitoring capability is also provided by the TBW designs as proposed here. Overall, the designs described here are consistent with the roles and capabilities attributed to the TBW in [1].

2 Design Concept

The first design we will consider uses exactly the same hardware proposed for the BFU as described in [4]. A TBW would nominally integrate into the DP1 subsystem as if it were a 4th BFU; although in fact it could take any position in the DP1 daisy chain. The same preliminary choice of FPGA – the Altera Cyclone III EP3C25 324-pin FBGA – is made here, and the interfaces are identical.

Each FPGA receives the 98 MSPS × 24-bit data from each of the 4 antennas (i.e., the four RF signals received from 2 stands). The data from each antenna is delivered to a dedicated asynchronous dual-clocked 24-bit wide first-in first-out (FIFO) buffer of length 4K (4096) samples, implemented on the FPGA. A fifth 16-bit × 4K-sample FIFO is used to collect signaling bits accompanying the data; with 4 bits per sample per antenna possible. Thus in this scheme $N_{bs} = 28$.

The normal mode of operation would be as follows: Once triggered, the TBW would immediately and coherently acquire data and signaling information, organized as described above. After 41 μs (the time it takes to fill a 4K buffer at 98 MSPS) the acquisition would automatically stop. Captured data would be offloaded from the TBW to the DAC by moving the contents of FIFO buffers, two at a time, through the paths identified as "LC partial beam" and "RC partial beam" in [4]. This transfer occurs at the same rate as the input data acquisition; therefore the offload time is 256 times longer than the capture time. Thus, this scheme achieves a best possible "duty cycle" of about 0.4%. Also, note that 4K captures imply best possible spectral resolution of about 24 kHz.

Alternative modes are possible by trading off N_{bs} for increased capture length. The extreme limit might be $N_{bs} = 4$ (2-bits "I" × 2-bits "Q", and no signaling bits). In this case the capture

length is 28K, corresponding to about 292 μs and a spectral resolution of 3.4 kHz.

It should be noted that these schemes consume only about 75% of the available memory and only about 1% of the available logic elements on the EP3C25. Thus excellent flexibility exists for implementing run-time processing as part of the capture. For an idea of what kind of processing is possible given these resources, see [2] and [4].

3 FPGA Implementation

To confirm the efficacy of the scheme proposed in the previous section, firmware was developed in Verilog HDL using Altera's Quartus II software, Ver. 7.2. An approach similar to that used in [4] was employed. For example, in order to reduce design time, "resync," "deformat," and "reformat" blocks were not implemented and LVDS pins were not utilized. Instead, all inputs and outputs were routed directly and automatically to available (single-ended) pins, which are available in sufficient quantity. This is not expected to impact conclusions here since ample LVDS channels are available and the logic required to implement the omitted functions is quite lean. FIFOs were, of course, implemented.

The $N_{bs} = 28$, length-4K design described in the previous section requires 75% of the available on-chip memory, about 1% of the available logic elements, and achieves timing closure for input clock rates up to 198 MHz; thus the desired 98 MHz is a comfortable fit.

4 Extending Capture Length

In order to extend capture length, it will be necessary to modify the TBW hardware, deviating from the design of BFU boards.

One approach might be to use a different FPGA, having additional memory. In order for this to be "interesting," the memory increase should be at least an order of magnitude. A review of what is available from Altera indicates that this is not currently an attractive option. The largest memory upgrade identified in a form factor that would permit reuse of the existing BFU design was for the Stratix II EP2S60, which has 2485 Kb (increase by a factor of ~ 4), a slightly larger form factor (23 mm $\times$ 23 mm vs. 19 mm $\times$ 19 mm), and a cost about 15 times greater ($760 ea.). Using this FPGA would increase the cost of a TBW by about $90K; hard to justify for a mere 4 times increase in capture length.

Another approach is to use the same FPGA with external RAM. Altera provides an complete example of interfacing to single data rate (SDR) synchronous dynamic RAM (SDRAM), including Verilog source and performance metrics, in [5]. Using the popular 128 Mb Micron MT48LC8M16A2 SDR SDRAM [6], the design example achieves a write throughput of 212 MB/s using 32-bit transfers. The impact on FPGA resources is very small; only about 200 logic elements (LEs) with $f_{max} =$ 100 MHz. Using this scheme, 4 SDRAM chips would be required (thus, about 800 LEs and four times as many pins). Each chip would support one antenna at $N_{bs} = 16$ (as opposed to the maximum value of 28) at 98 MSPS, which is 196 MB/s. The chip capacity of 128 Mb/s at $N_{bs} = 16$ is $8 \times 1024 \times 1024$ samples (again at 98 MSPS), which corresponds to a capture length of 85.6 ms and a best-possible spectral resolution of 12 Hz. This part is available in a 54-ball VFBGA form factor, which has a footprint 8 mm $\times$ 7 mm, so 4 of these could conceivably be integrated into the existing BFU-based circuit board, especially if the connectors for continuing the DP1 daisy chain were removed. This would require that the TBW be the last subsystem in the DP1 daisy chain. Each SDRAM chip currently costs in the range $5–$10, so the worst case additional cost would be about 128 $\times$ 4 $\times$ $10 = $5120 per TBW.

5 Issues to Address in Future Versions of this Document

1. Many issues are the same as that for the BFU since the proposed hardware implementations are very similar (perhaps identical); please consult BFU preliminary design documents.

2. Repeat Cyclone III synthesis including the actual interfacing and variable-N_{bs} functionality, LVDS pins, etc.

6 Document History

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

References

[1] S. Ellingson, "LWA Station Architecture Ver 0.6," October 9, 2007.

[2] S. Ellingson, "ADC Sample Rate and Preliminary Design for a Full-RF ADC Post-Processor," LWA Memo 101, September 11, 2007.

[3] S. Ellingson, "DP1 Daisy Chain Preliminary Interface Control Document Ver. 0.1," Sep 18, 2007. Online: http://www.ece.vt.edu/swe/lwav/.

[4] S.W. Ellingson, "BFU & DP1 Preliminary Design," LWA Memo 108, November 4, 2007.

[5] Altera, Inc., "SDR SDRAM Controller" (White Paper), Ver. 1.1, August 2002.

[6] Micron, Inc., "SDRAM MT48LC8M16A2" (data sheet), 10/07. http://www.micron.com.