

FP7 - Grant Agreement no. 283393 - Radionet3

Project name: Advanced Radio Astronomy in Europe
Funding scheme: Combination of CP & CSA
Start date: 01 January 2012
Duration: 48 months



Deliverable 8.15

Uniboard² - Effectiveness of Green Measures:
Digital Receiver

Due date of deliverable: 30 Nov 2015
Actual date of deliverable: 30 Nov 2015
Deliverable Leading Partner: INAF



An European project supported within the 7th framework programme



1 Document information

Document name: Report on effectiveness of green measures: Digital Receiver
 Type Document
 Revision 1.0
 WP 8
 Authors Giovanni Comoretto
 INAF report 3/2015

1.1 Dissemination level

Dissemination Level		
PU	Public	X
PP	Restricted to other programme participants (including the Commission Services)	
RE	Restricted to a group specified by the Consortium (including the Commission Services)	
CO	Confidential, only for members of the Consortium (including the Commission Services)	

1.2 Document history

Revision	Date	Author	Modification/Change
0.1	2015-08-21	Comoretto	Initial draft
0.8	2015-11-25	Comoretto	Revised complete version

1.3 Distribution list

ASTRON Andre Gunst, Eric Kooistra, Sjouke Zwie, Danil van der Schuur, Harm Jan Pepping

JIVE Arpad Szomoru, Jonathan Hargreaves, Salvatore Pirruccio, Sergei Pogrebenko, Paul Boven, Harro Verkouter

UMAN Ben Stappers, Prabu Thiagaraj

INAF Gianni Comoretto

BORD Benjamin Quartier, Alain Baudry, Stephane Gauffre

UORL Cedric Dumez-Viou, Rodolphe Weber, Nicolas Grespier

MPG Guenter Knittel, Gundolf Wieching

1.4 Terminology

10GbE: 10 gigabit per second Ethernet

40GbE: 40 gigabit per second Ethernet

ADC: Analog to Digital Converter

bps: Bits per second

BW: BandWidth

DDR: Double Data Rate. DDR3 and DDR4 refer to the 3rd and 4th version of the standard, respectively

DSP: Digital Signal Processing

Firmware: Embedded or real-time code that runs on a microprocessor (e.g. written in C), or describes a programmable logic (e.g. written in HDL)

FFT: Fast Fourier Transform

FPGA: Field Programmable Gate Array

Hardware: Boards, sub-racks and COTS equipment

HDL: Hardware Description Language

HMC: Hybrid Memory Cube: high density memory standard, with high speed serial interface

IO: Input-Output

IP: Intellectual Property

JESD204: A JEDEC (Joint Electron Device Engineering Council) standard for serial ADC interface

JTAG: Joint Test Action Group standard for boundary scan testing and programming of programmable devices

LAN: Local Area Network

PHY: physical interface (layer 1 of OSI model)

QSFP+: SFP for 40Gb Ethernet

PFB: Polyphase Filterbank

RFI: Radio Frequency Interference

SFP: Small Form-factor Pluggable transceiver. An optical interface module, available with different performance and range characteristics, pluggable in a common socketed cage.

SFP+: SFP for 10Gb Ethernet

TCP: Transmission Control Protocol, standard safe Internet protocol

UDP: User Datagram Protocol, standard stateless, low overhead Internet protocol

WAN: Wide Area Network

1.5 References

References

- [1] L. D'Addario: "How to Implement SKA Digital Signal Processing So That It Uses Very Little Power" Workshop on Power Challenges of Mega-Science, Moura, Portugal., 20 June 2012

- [2] Hargreaves, J.: “Report on effectiveness of green measures: correlator”, UniBoard2 Design Document Deliverable 8.14, (2015).
- [3] Comoretto, G.: “Revised Firmware Design Document Uniboard2 Digital Receiver”, UniBoard2 Design Document Deliverable 8.11, (2015).
- [4] Analog Devices: “JESD204B Survival Guide”
http://www.analog.com/static/imported-files/tech_articles/JESD204B-Survival-Guide.pdf
- [5] G. Comoretto, A. Russo, G. Tuccari, A. Baudry, P. Camino, B. Quertier: “Uniboard Digital Receiver Design document”, Arcetri Technical Report 5-2011
http://www.arcetri.astro.it/images/data/Reports/11/5_2011.pdf
- [6] Manley, J.R.; “A Scalable Packetised Radio Astronomy Imager”, PhD Thesis. Department of Electrical Engineering, University of Cape Town, 2014.
- [7] P. Camino, D. Dallet, B. Quertier, A. Baudry, G. Comoretto, B. Le Gal: ”A Decimation Filter for Very Large Band Signal in Radioastronomy” PRIME 2007 15.3
- [8] Arpad Szomoru: “UniBoard2 Work Package description”, RadioNet3 283393 (2011)
- [9] Gijs Schoonderbeek: “UniBoard² Architecture - Hardware design document”, Astron report INFRA-2011-1.1.21 (2014)

Contents

1	Document information	2
1.1	Dissemination level	2
1.2	Document history	2
1.3	Distribution list	2
1.4	Terminology	2
1.5	References	3
2	Introduction	6
2.1	Background	6
2.2	Typical FPGA power consumption	6
2.3	Application specific considerations	7
3	Remedies	7
3.1	FPGA technology used	7
3.2	Automatic optimization	8
3.3	Standby modes	8
3.4	Architecture consideration	9
3.5	Algorithm optimization	9
3.6	DC offset schemes	9
4	Conclusions	10

2 Introduction

2.1 Background

Digital processing for radioastronomical applications requires both a large bandwidth, as the instrument sensitivity is proportional to the signal bandwidth, and a complex data processing, including channelization, filtering, RFI excision. These requirements reflect in a significant power consumption. The typical environment for a digital receiver is at a telescope site, where power can be limited, and expensive, with an energy cost during one year operation that is typically comparable to the capital cost of the hardware. For example, early estimates of the power needed for signal processing for the SKA project range from 29 MW to 85 MW[1].

These figures include a 47% overhead for power supply loss and cooling. They translate to an annual cost of between 30 and 89 million euros, based on a price of E0.12 per kWh, a significant proportion of the SKA operating budget.

Therefore particular care is required for minimizing the power used by the digital hardware. Most of these techniques are common to all digital systems, and the analysis described in [2] applies also to this application.

2.2 Typical FPGA power consumption

As any CMOS device, FPGAs use power both statically (i.e. independently from the performed function) and dynamically. Dynamic power occurs during power switching, both for resistive losses during the commutation, and because of capacitive loading. Static power is typically due to tunneling in the very small transistors, and to resistive loads in transmission lines. Portion of the device can be almost completely turned off, reducing total power consumption to a typically smaller *standby power*.

Power behaviour is different for different portions of the FPGA:

- Core logic. Core logic performs the bulk of the logic processing, and exploits the ever increasing level of integration. As transistor size decreases, static power, dominated by transistor leakage, increases due to the increasing relevance of tunnel effects. This is in part mitigated by decreasing the operating voltage. Dynamic power decreases for smaller transistor sizes, due to reduction in both gate capacity and operating voltage.
- Input/output. The external signal voltage cannot decrease as the internal core voltage, due to susceptibility to interfering signals. Signal lines must be operated at a low impedance, and are typically terminated with a resistive load, resulting in significant driving currents, both static and dynamic. For Uniboard² the largest external interface is represented by the two DDR4 memory banks, with about 800 mW required by each bank interface.
- High speed transceivers. These lines must reliably operate at speeds in excess of 10 Gbit/s, requiring low impedances, high driving currents, and complex fast driving electronics. A single 10GbE link requires about 225 mW, not including the power in the optical transceiver and QSFP+ cage. An HMC memory module uses 16 links at 10 Gb/s, requiring 3.6 W, i.e. also 225 mW per link.

The Uniboard FPGA digital receiver application[3] requires very little power for input/output, as the only significant I/O peripheral (DDR memory) is not used. Core logic is heavily used, with about 50% of the sparse logic and nearly 100% of the hard multipliers involved in the design, and runs at near the maximum operating frequency. Transceivers are not heavily used, as the data is to be transmitted over a commercial link that is the true system bottleneck. For example a 2 GHz bandwidth, with 2 bit quantization, requires 8 Gbit/s, i.e. a single 10Gb transceiver. ADC converters are connected using serial links, either with packetized protocols

over standard UDP, or using the JESD204 protocol[4]. A total of 4 10Gb links are required for the assumed 4 bit, 8 GSample/s ADC.

2.3 Application specific considerations

The digital receiver is not a particularly power intensive application. It is conceived as a small component of the telescope instrumentation, and is implemented as a single board. The Uniboard version[5] processes a total bandwidth of 4 GHz, either as a single signal or as two signals of 2 GHz each, and the resulting processed data are presented as a total of over 4 10GbE channels. The Uniboard² version considered here is implemented in a single FPGA, i.e. 1/4 of a board. Considering 4 independent channels, for a total processed band of 32 GHz, the total power consumption is less than 100 W. Power saving is thus minimal for the intended application, but may become relevant if the design is used in an interferometric environment, with hundreds of individual antennas.

3 Remedies

3.1 FPGA technology used

The FPGA technology is quickly evolving. The first version of the digital receiver application[5], running on Stratix5 Altera FPGAs (60 nm process), required a total of 45 W to operate the Digital Receiver application, on 8 FPGAs. A large fraction of this power, 12 W, was used for chip-to-chip communication. The same design, on a Arria10 based system (20 nm, engineering sample 1), requires 18 W, on a single FPGA.

A comparable improvement is expected going to the 14 nm Stratix10 technology. An important power saving strategy consists then in writing the application, and possibly design the hardware, in order to quick exploit any change of technology.

Any change of hardware implies an energy waste due to the embedded energy for the hardware construction, that in first approximation is proportional to its cost. From a purely economic point of view, a detailed analysis [6] shows that switching to a new, more efficient technology becomes advantageous on the average after a 3 year period. As an example, figure 1 shows both the total cost (ownership plus power) for the MeerKat correlator using the latest available technology, up to the assumed end of life in 2050, and the power cost needed to operate the instrument for the remaining time for successive generations of FPGAs. When the former drops below the latter, a technology upgrade becomes advantageous.

This fits well with the experience of Uniboard to Uniboard 2 transition, and with the expected time for the Stratix10 version of the Uniboard 2. In this latter case, the pin-to-pin compatibility of the two FPGAs allow a very easy and smooth transition, without all the costs derived from the development of a new board. A significant power saving, of the order of 15%, can be achieved just by switching from the pre-production to the production versions of the FPGA.

A more aggressive approach consists in switching to an ASIC version of the application. This allows for a tenfold decrease in logic power consumption, but at the expenses of flexibility. FPGA vendors provide a roadmap for deriving an ASIC version form a frozen design, simply by substituting the FPGA programming memory and switches with a metallization mask. The derived ASIC is not optimized, but this solution allows for an initial testing of the design using programmable logic, and a final production version using ASIC.

Past experience shows however that design modifications, due to subtle digital processing effects unaccounted in the initial design, design improvements, or engineering change requests, may occur even after the initial commissioning of the instrument. This seriously limits the usefulness of this approach in the radioastronomic domain.

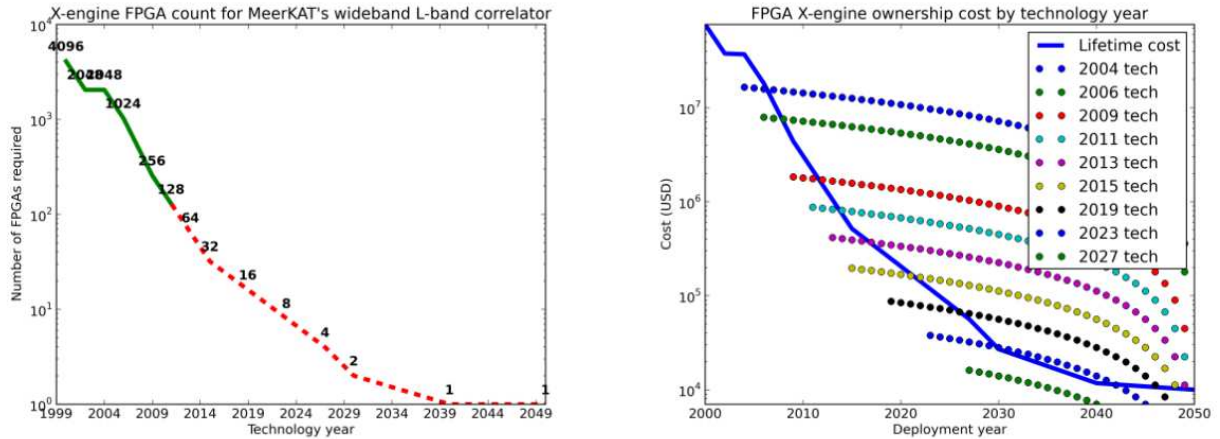


Figure 1: Left: number of FPGAs needed for the MeerKat correlator, as function of available technology, projected up to 2036. Right: continuous line is total cost of ownership using last available technology (purchase plus power), dotted lines are operating cost to end of life of the instrument. From [6]

Category	Default settings	Extra effort synthesis	Percent change	Extra effort synth. & fitter	Percent change
Device Static	3.609	3.612	0.08	3.580	-0.81
Core Dynamic	4.068	4.105	0.91	3.894	-4.28
Total Power Dissipation	15.175	15.216	0.27	14.972	-1.33
Compile Time (hh:mm)	2:02	2:08		3:00	

Table 1: Effect of the Powerplay Extra Effort setting in Quartus: UniBoard2 test design (from [2])

3.2 Automatic optimization

The Quartus synthesis software can be configured for optimization of speed, power, or both. The optimization is performed either in the synthesis phase (translation of abstract functionality to FPGA logics), or both in the synthesis and fitting phase (placing the logics inside the FPGA fabric). Tests performed using a sample design shows that the achievable power reduction is in the order of a few percent[2].

The method has the advantage of being completely automatic, and costs just some more computing time in the final iteration of the design process. The power gain is however marginal.

3.3 Standby modes

Dynamic power is cut to virtually zero by disabling the clock in the unused portions of the design. A good design practice is to provide a constant input to any part of the design, when not in use, but the internal clock network contributes to about 5-10% of the total power, even when the logic itself is not switching. A better approach is to partition the clock network and use controllable clock buffers, that are actively disabled when not needed. The digital receiver

Representation size	8 bits	10 bits	12 bits	18 bits
Natural word size	19%	15%	12%	8%
18 bits word size	41%	35%	28%	8%

Table 2: Power reduction in applying an offset to a Gaussian noise, for different word size

design is composed of 64 identical units, of which typically only a portion is used at any time. As it is difficult to provide 64 independent clock networks, the units are grouped in larger groups, that are enabled or disabled as a whole.

3.4 Architecture consideration

A large fraction of the power used by a FPGA board comes from data transport across FPGAs. For a typical design using most of the computing resources and 24 transceivers (240 Gbit/s, i.e. about 15 GHz of RF bandwidth equivalent), the transceivers use about 40% of the available power. If an optical transceiver is used, this adds a comparable amount of power. Minimization of the number of links is then mandatory.

Using large FPGAs, as in Uniboard² compared to the original Uniboard, the number of inter-chip connections is drastically reduced. The original filterbank application used 8 FPGAs, with a total of 32 interconnecting links using about 10 W. The Uniboard² version of the application fits in just a single FPGA, drastically reducing the amount of interconnections.

3.5 Algorithm optimization

Filter design can be optimized for reduced power consumption. An example of this approach is the broadband filter used in the ALMA channelizer[7]. The original design, using a dual stage filter, uses about 110 W for a 32 channel digital receiver. Adopting a different design, with a 3 stage filter, reduced the power by 20%, to 80 W per board. These optimizations have already been considered in the original Uniboard design, so no further improvement was possible.

3.6 DC offset schemes

Consumption in the digital core is dominated by dynamic power, that is proportional to the number of transitions in all signals.

When a signed value changes sign, i.e. when the signal crosses the zero, all most significant bits change value together. A typical radioastronomic signal has a Gaussian distribution with zero mean, and a spectrum using most of the available digital bandwidth, producing a zero crossing about 50% of the time. The signal RMS amplitude is typically 1/10 of the available range, which means that the 3-4 most significant bits are usually identical. Reducing the frequency of the zero crossing can significantly reduce the dynamic power consumption.

Adding a constant offset, of the order of the RMS of the digital value, together with a small decrease of the signal amplitude, reduces the number of zero crossings by about 60% against a 20% decrease of the dynamic range. The total decrease in single bit switching is typically much less, as it affects only 3 bits of the signal representation, and becomes less significant as the number of bits in the signal representation increases (table 2). It can be however significant when hard-core multipliers, that have a fixed size of typically 18×18 bits, are used to represent small integers, of 8-10 bits.

The power reduction applies only in systems that preserve a DC offset, like low pass filters. It is ineffective in frequency translation applications, like FFT or digital mixers. But as the digital filtering, in digital baseband converters and polyphase filters, represent a large fraction

of the processing and is typically performed on small amplitude (few bits) signals using 18 bit multipliers, the power saving can be significant, up to 20% of the core dynamic power, or about 6% of the total.

4 Conclusions

Reducing the power consumption of a digital circuit is not an easy task. The most effective approach relies in the use of larger and less power hungry integrated circuit processes. Typical power saving of 50% is achieved in the transition to a new CMOS process. This must come in synergy with an effort in reducing the number of external interconnections, that are less affected by technological advances, and in a standardization in the design guidelines that ease the transition to the new technologies, when they become available.

Other approaches may contribute with small power savings. Those include:

- Better synthesis and placing algorithms, at the expense of longer compilation times
- Careful architectural design, focused on power minimization
- Partitioning of the design into smaller functional groups, that can be individually disabled when not used
- Reduced number of zero transitions in the signal

The power gain in any of these measures is of the order of 1-8%, and cumulatively they contribute to as much as 20%.