

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 month



Deliverable 8.14

**Report on effectiveness of green measures:
correlator**

Due date of deliverable: 2015-08-31

Actual submission date: 2015-09-02

Deliverable Leading Partner: JOINT INSTITUTE FOR V.L.B.I. IN EUROPE (J.I.V.E.)

1 Document information

Document name: UniBoard2 – Effectiveness of Green Measures: Correlator

Type Report

WP 8 (UniBoard²)

Authors J. Hargreaves, (JIVE)

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)	

UniBoard² Effectiveness of Green Measures: Correlator

Deliverable 8.14
J Hargreaves, JIVE
25th August 2015

1.1 Terminology

DDR3,DDR4	Double Data Rate memory interface standards
DSP	Digital Signal Processing
ES	Engineering Sample
EVN	European VLBI Network
FPGA	Field Programmable Gate Array
HMC	Hybrid Memory Cube
IO	Input/output
PHY	The physical layer of an interface
MAC	Media access controller, part of the data link layer of an interface
MTU	Maximum Transmission Unit in a packet switched network
NIOS	An embedded processor used in Altera FPGAs
SFP+	Small Form factor Pluggable interface
QSFP	Quad Small Form Factor Pluggable Interface
SKA	Square Kilometer Array
MSPS	Million Samples Per Second
GSPS	Giga Samples Per Second
UNB	UniBoard
VHDL	Very high-speed integrated circuit Hardware Description Language
1GbE, 10GbE	One- or ten- gigabit per second Ethernet

1.2 References

1. "How to Implement SKA Digital Signal Processing So That It Uses Very Little Power" D'Addario, L., Workshop on Power Challenges of Mega-Science, Moura, Portugal., 20 June 2012
2. "UniBoard² Design Document: A Low Spectral Resolution EVN Correlator for Continuum Processing" UniBoard² Project Deliverable 8.11, Hargreaves, J., 25 June 2015
3. "Leveraging HyperFlex Architecture in Stratix 10 Devices to Achieve Maximum Power Reduction", Won, M., Altera white paper WP-01253-1.0, June 2015
https://www.altera.com/content/dam/altera-www/global/en_US/pdfs/literature/wp/wp-01253-leveraging-stratix10-hyperflex-for-maximum-power.pdf

1.3 Table of Contents

1.1	Terminology	1
1.2	References	1
1.3	Table of Contents	2
2	Introduction	3
2.1	Background	3
2.2	Typical FPGA Power Consumption	3
3	Discussion and Remedies	5
3.1	Quartus Optimizations	5
3.2	Production versus Engineering Samples	6
3.3	Newer FPGAs	7
3.4	Minimize IO	7
3.5	Standby Modes	8
3.6	Algorithm Optimization	8

2 Introduction

2.1 Background

In general, radio astronomy receivers are constructed in remote areas to minimize radio interference from human habitation, but this means that electrical power supplies may be limited or expensive.

Significant power is required to digitize, transport, store and process data. In addition to power consumed in the FPGAs, CPUs and ASICs doing the computation, power supplies are typically 85% efficient. Further power is required for the fans and air conditioning systems to remove heat from the signal processing circuitry.

For example, early estimates of the power needed for signal processing for the SKA project range from 29MW to 85MW [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.

This document explores ways of reducing the power consumption of FPGA signal processing designs implemented on UniBoard and UniBoard².

2.2 Typical FPGA Power Consumption

Altera's Quartus design software includes a 'Powerplay' power analysis tool to estimate power consumption within an FPGA for a given compiled design. The tool reports power usage broken down into several categories:

Static, Standby and Dynamic

Static power results from the underlying leakage currents of the entire device. At constant temperature static power is independent of clock speeds and device utilization. Note that in practice, as dynamic or IO power increases, static power will also increase because the operating temperature of the device will rise.

Dynamic power is dissipated whenever signals change state, so includes the power needed to process and transport signals and clocks within the FPGA.

For IO and transceiver blocks, 'Standby' power accounts for the additional bias currents needed when the block is switched on in the design, but idle.

Core, IO and Transceiver

Transceivers and general purpose IOs are driven by several different power supplies and voltages, and so are reported individually. The core includes the multipliers, logic, memory and routing which comprise the fabric of the FPGA.

Logic and Routing

Within the core, power used in logic elements can be distinguished from power used to move signals around the IC. The power used in different parts of the design hierarchy can be seen, as can the power used in different types of physical structures (DSP, memory and so on).

The following table summarizes the Powerplay power estimates for one of the test designs used to verify the UniBoard² hardware. This design comprises twenty-four 10GbE interfaces including transceiver PHYs and MACs, and a NIOS controller. The design was compiled for the Arria 10 engineering sample (part number 10AX115U4F45I3SGES) installed on the first prototype UniBoard² board. The design occupies approximately 20% of the FPGA logic, 14% of the memory, and 10% of the registers, but no DSP resources.

Category	Power in milliwatts	Percent of Total
Device Static	3609	24
Core Dynamic	4068	27
IO Standby	108	1
IO Dynamic	72	<0.5%
Transceiver Standby	1661	11
Transceiver Dynamic	5658	37
Total Power Dissipation	15175	100

Table 1: Summary power estimates for 24-transceiver UniBoard² test design

Static power is nearly a quarter of the total in this relatively sparse design. Static power can be minimised by minimising the number of FPGAs required in the system. This implies maximising resource utilization and clock rates within each FPGA. It should be noted however that timing bottlenecks are more likely in designs using more than 85% of an FPGA at clock rates higher than 250MHz. This can result in longer design times due to the need to add pipelining registers or manually guide the place and route tools.

The IO power consumption is negligible in this design, since very little general purpose IO is used: a 1GbE port for control, a clock input, and some GPIO pins to drive status LEDs.

The transceivers by contrast draw nearly half the power, perhaps not surprising as they provide a total bandwidth of 240MHz full duplex, equivalent to the entire transceiver throughput of UniBoard¹. IO and transceiver power will be discussed further in Section 3.4.

Accuracy of Estimates

Dynamic power is sensitive to the toggle rates of the signals in the design. More accurate power estimates can be obtained by simulating the design to estimate toggle rates. In this document a default 12.5% toggle rate was used for all input signals, since simulated data was not available for all designs. This allows comparisons to be made, but gives only a rough guide to the true power consumption.

As a guide, the power dissipation measured on the UniBoard² hardware for the design in table 1 was 30% higher than the estimates. The difference may be explained by higher than estimated toggle rates, higher ambient temperature, and external devices sharing the FPGA power supplies.

3 Discussion and Remedies

3.1 Quartus Optimizations

The Altera Quartus software can be set to optimize the design for power consumption at both the synthesis and fitter (place and route) stages. The level of effort is set to 'Normal' by default, but can be set to 'Off' or 'Extra Effort'. This test compares the 'Normal' setting with 'Extra Effort' for synthesis only, and 'Extra Effort' for both synthesis and fitter.

Table 2 shows the results for the same test design used in table 1. The IO and transceiver figures are omitted because they are not affected by the optimization.

Category	Power (mW) default settings	Power (mW) Extra effort synthesis only	Percent change from default	Power (mW) extra effort synthesis and fitter	Percent change from default
Device Static	3609	3612	0.08	3580	-0.81
Core Dynamic	4068	4105	0.91	3894	-4.28
Total Power Dissipation	15175	15216	0.27	14972	-1.33
Compile Time (hh:mm)	2:02	2:08	-	3:00	-
Timing slack (ps)	-709/-433	-940/-456	-	-978/-433	-

Table 2: Effect of the Powerplay Extra Effort setting in Quartus: UniBoard² test design

FPGA compilation is a compromise between fitting the design in the chip, meeting the timing constraints, power efficiency, and a reasonable compile time. Table 2 shows two additional parameters: the compile time and the worst-case timing slack for the two main clock domains in the design. The negative timing numbers mean that this test design did not meet timing – more negative is worse.

In this case switching on 'extra effort' at the synthesis stage had little effect on the power dissipation. The core dynamic figure is slightly worse, but within the random difference that could be expected between successive runs. The compile time increased marginally as did the timing error.

When the 'extra effort' option was switched on for the fitter stage as well as synthesis, the dynamic power improved by 4%, though the compile time was 50% longer. The timing result was not significantly worse than the previous, synthesis-only, case.

The comparison was repeated using a mature UniBoard1 design, the correlator engine for the JIVE UniBoard correlator (JUC). This design has a system clock frequency of 260MHz, logic resource utilisation of 30% and contains 532 18x18bit multipliers. The results are shown in Table 3.

Category	Power (mW) default settings	Power (mW) Extra effort synthesis only	Per cent change from default	Power (mW) extra effort synthesis and fitter	Per cent change from default
Device Static	1810	1808	0.11	*	-
Core Dynamic	6568	6550	0.27	*	-
Total Power Dissipation	9737	9716	0.22	*	-
Compile Time (hh:mm)	2:06	2:23	-	*	-
Timing slack (ps)	85/186	135/205	-	*	-

Table 3: Effect of the Powerplay Extra Effort setting in Quartus: JUC 'X' node

Again in this case the improvement in power dissipation due to the optimization was small, though there was a slight improvement in timing margin. When the fitter 'extra effort' option was enabled, the design failed during place and route due to conflicts with location constraints. However the location constraints were put in early in the design cycle to help timing closure. The conclusion is that there is a trade off between optimisation for timing performance and power efficiency during place and route. The power optimisations would have to be enabled early in the design cycle, before adding location constraints, to avoid the conflict.

3.2 Production versus Engineering Samples

Production devices can be expected to have lower leakage currents, and hence lower power consumption than the figures for engineering samples shown in tables 1 and 2. In the following table power estimates for the ES device are compared to the production device (part no 10AX115U4F45E3SG) to be installed on the first batch of production UniBoard²s.

For this comparison a pinning design was used. This design uses relatively little of the core resources of the FPGA, but all the transceivers and IO pins are used, including two 72-bit DDR4 controllers.

Category	ES Device Power (mW)	Production Device Power (mW)	Per cent Change
Device Static	7851	5556	-29
Core Dynamic	539	454	-16
IO Standby	3701	3701	0
IO Dynamic	1091	1385	27
Transceiver Standby	6381	6643	4
Transceiver Dynamic	21037	16569	-21
Total Power Dissipation	40601	34308	-15

Table 4: Comparison of dissipation in ES and production Arria 10 devices

The production device operates at a core voltage of 0.9V, against 0.95V for the ES device. The power consumption of the production device is mostly lower, except in the IO dynamic category. This appears to be due to an increase in the power consumed by the 'IO Digital' sub-category from 845mW to 1137mW. This category includes the two DDR4 memory controllers, and may possibly be explained by the fact that DDR controller hardware in the production device will offer greater functionality and performance than that in the ES device.

3.3 Newer FPGAs

Newer FPGAs can be expected to consume less energy per unit of computation due to lower core voltages and smaller line widths: smaller gates mean less charge has to move each time a signal changes state. Comparing the first line in Tables 2 and 3, it can be seen that the static power consumption of the 20nm Arria 10 FPGA is about twice that of the 40nm Stratix IV, even though the Arria 10 device has roughly four times the computational power [2].

Further power improvements can be expected for the Stratix 10 devices resulting from the move to 14nm technology, higher clock speeds obtained through the 'HyperFlex' architecture, the ability to switch off unused logic blocks, and a further reduction in core voltage to 0.8-0.85V [3].

3.4 Minimize IO

It can be seen in table 4 that power dissipation transceivers and IO is significant: 82% of the total in this case, made up of 30% standby and 52% dynamic. These figures are for an IO pinning check design with negligible core resource usage, but still indicate that there is a high fixed power cost for each active interface. Some additional estimates of interface power consumption are given below. The figures are for the 10AX115U4F45E3SG production device. In each case the static power is greyed out, and omitted from the totals, because it is independent of whether the interface is active in the design.

Hybrid Memory Cube (HMC) interface (16 transceiver lane)

A design containing two instantiations of Altera's HMC interface example design was created. Each interface consists of a 16 lane bi-directional transceiver PHY, plus driver logic including a test pattern generator/checker. After running 'Powerplay' the following estimates of the power per HMC interface were derived:

Category	Power in milliwatts
Core dynamic power for the test pattern and driver logic	335
Transceiver static power for the 16 lane PHY	41
Transceiver standby power for the 16 lane PHY	812
Transceiver dynamic power for the 16 lane PHY	2456
Total per HMC interface	3603

Table 5: Dissipation estimates for an HMC interface

DDR4 interface (72 bit)

The figures are derived from the pinning design and include power for the hard DDR4 controller and the PHY.

Category	Power in milliwatts
Core dynamic power DDR4 hard controller	409
Core standby power DDR4 hard controller	150
Core routing power DDR4 hard controller	9
IO static power for the DDR4 PHY	0.2
IO standby power for the DDR4 PHY	235
IO dynamic power for the DDR4 PHY	11
Total per DDR4 interface	814

Table 6: Dissipation estimates for a DDR4 interface

Single 10GbE transceiver

Category	Power in milliwatts
Transceiver static power	6
Transceiver standby power	68
Transceiver dynamic power	157
Total per 10GbE interface	225

Table 7: Dissipation estimates for a single 10GbE PHY

The above estimates do not include the dissipation in the HMC and DDR4 modules, and the QSFP cages in the case of optical board-to-board communication. Power consumption for IO can be minimized by:

- Choosing an architecture that relies less on external dynamic memory and more on internal FPGA memory. An example is the proposed low spectral resolution EVN correlator [2] that saves power by eliminating an external memory based corner-turning stage.
- Packing the design into the smallest number of large FPGAs available, to reduce the need to transfer data between FPGAs. The proposed Stratix 10 version in [2] combines the 'F' and 'X' parts of the correlator into one FPGA, eliminating the need to send intermediate data between FPGAs on the ring mesh.
- Use 10GbE ports full duplex where possible to reduce the number of ports active, and hence the total standby power.

3.5 Standby Modes

Further operating power reductions can be achieved by implementing standby modes. In the correlator example, the standby mode could be used between scans when there is no need to retain data or state.

Dynamic power falls automatically when no data processing is taking place, but it would also be possible to switch off unused clock networks. In the UniBoard² test design of Table 1, clock networks account for 825mW of the dynamic power. Altera FPGAs include a clock control/driver component called ALT_CLKCTRL. This provides an enable input that could be used to shut down all clocks except for those needed for the 1GbE and control circuitry.

In Stratix 10 devices it will be possible to go further and switch off unused logic blocs thus reducing the total static dissipation in standby mode.

DDR4 modules have a low power, non-data retention, standby mode claimed to reduce consumption by around 40%.

3.6 Algorithm Optimization

There is potential to reduce dynamic power consumption by optimizing the algorithm used to perform the calculation. Dynamic power is consumed whenever signals change state, so the goal of optimization is to reduce the number of bits changing state over time. As a trivial example, consider a state machine with three states IDLE, S0 and S1. Suppose it toggles frequently between S0 and S1 but rarely enters the IDLE state. In that case the second encoding shown below would be more energy efficient than the first because only one bit changes between S0 and S1.

State	Encoding 1	Encoding 2
IDLE	00	10
S0	01	00
S1	10	01

Table 8: Example of state machine encoding for energy efficiency

In some cases a signal does not need to be updated on every clock cycle. Typically the signal will be frozen, or latched, for several clocks using a flip-flop and clock enable signal as shown in Figure 1 below. The clock enable might only be active high every second, fourth or more clock but if the inputs are not similarly latched, the combinational function will be recalculated every clock, wasting power. It would be more energy efficient to latch the flip-flops driving the inputs where possible.

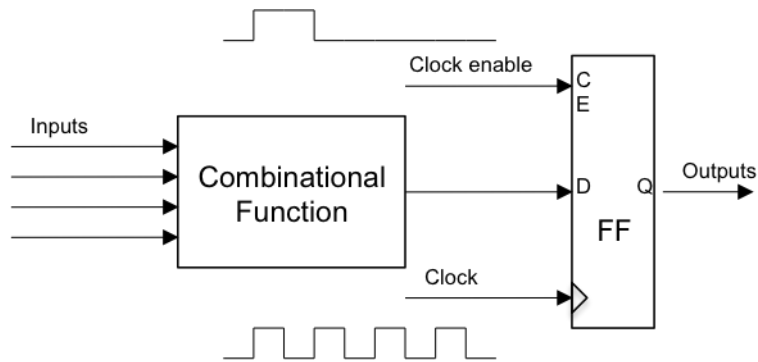


Figure 1: Gating signals to reduce power consumption

Truncation or rounding can be introduced to reduce signal path widths, provided the dynamic range remains within specification. This is particularly useful when the signals are encoded as two's complement signed integers. The encoding for the eight-bit case is shown below

Decimal Value	2's complement signed coding
127	01111111
1	00000001
0	00000000
-1	11111111
-128	10000000

Table 9: Two's complement signed integers

If the signal path is unnecessarily sign-extended, meaning extended to the left with copies of the sign bit, all the extra bits will wastefully change state every time the signal crosses zero. Depending on the signal statistics, an alternative encoding scheme might be more efficient. For example if the signal contains predominantly small values either side of zero, the Gray code shown in table 8 might be more efficient, though the data would normally have to be re-converted for standard DSP processing.

Decimal Value	2's complement signed	Gray code
7	0111	0100
2	0010	0011
1	0001	0001
0	0000	0000
-1	1111	1000
-2	1110	1001
-8	1000	1100

Table 10: Four-bit Gray encoding scheme

Other energy saving coding measures are listed below. In general simulation, Powerplay analysis or manual calculation would be needed to determine the most efficient option for a given design specification and clock rate.

- A function where an intermediate result is re-used in a calculation can be made more energy efficient by calculating once and storing, though at a cost of coding complexity.
- Functions such as multiplier-adders, typically found in filter structures, may be re-ordered to optimize energy efficiency.
- Look up tables may be more efficient than DSP resources for implementing multipliers with small bit widths.
- Invalid data should be forced to zero as early as possible in the signal chain.