Experience powering Xilinx Virtex-7 FPGAs

2013 JINST 8 C12037

(http://iopscience.iop.org/1748-0221/8/12/C12037)

PUBLISHED BY IOP PUBLISHING FOR SISSA MEDIALAB

RECEIVED: *November 11, 2013*
ACCEPTED: *December 11, 2013*
PUBLISHED: *December 20, 2013*

TOPICAL WORKSHOP ON ELECTRONICS FOR PARTICLE PHYSICS 2013,
23–27 SEPTEMBER 2013,
PERUGIA, ITALY

2013 JINST 8 C12037

Experience powering Xilinx Virtex-7 FPGAs

G. ILES,[a] **J. Jones,**[b] **and A. Rose**[a,1]

[a] *Imperial College, London, Blackett Laboratory,*
Prince Consort Road, London, SW7 2BW, U.K.
[b] *Iceberg Technology, Helland Mill,*
Mabe Burnthouse, Penryn, Cornwall, TR10 9JQ, U.K.

E-mail: awr01@imperial.ac.uk

ABSTRACT: The MP7 processor card is an all-optical signal processor based on mid-range Xilinx Virtex-7 FPGAs (XC7VX485T or XC7VX690T) in a 1927-pin package. The current version of the card uses the large serial bandwidth of these FPGAs to provide a 0.75+0.75Tbps bidirectional interface, although changes are currently being implemented to push the bandwidth to 0.94+0.94Tbps. The MP7 card was designed before Virtex-7 Engineering Silicon was available and so a lot of design work relied on preliminary documentation and discussion with Xilinx engineers. The MP7 card was, therefore, the first card in the CMS experiment to use Xilinx 7-series FPGAs. The experience of designing for and powering the Virtex-7 are discussed, along with lessons learned. The challenges associated with cooling an FPGA dissipating of-the-order 40W is also discussed.

KEYWORDS: Digital signal processing (DSP); Trigger concepts and systems (hardware and software); Modular electronics; Digital electronic circuits

[1] Corresponding author.

doi:10.1088/1748-0221/8/12/C12037

Contents

2013 JINST 8 C12037

1 The MP7 processor card and the Xilinx Virtex-7 FPGA

The MP7 processor card [1, 2] (see figure 1) is a 0.75+0.75Tbps all-optical signal processor based on mid-range, Xilinx Virtex-7 FPGAs [3]. The MP7 has been designed to be compatible with either the XC7VX485T or XC7VX690T FPGAs, which come in pin-compatible FFG1927 packages, differing in the amount of logic and the number of available serial links. The pin-out for the FFG1927 package is shown in figure 2, with the pins dedicated to powering highlighted for clarity. It is striking that, of the 1927 pins, approximately half are dedicated to powering and grounding the chip.

Unlike previous generations of FPGA, or even the Kintex or Artix members of the Xilinx 7-series, the Virtex-7 includes a design-choice which reduces the number of different voltages required: all pins on the Virtex-7 are of the "High-Performance" variety and so do not support I/O above 1.8V. This choice does, however, lead to its own problems, necessitating level-translation when communicating with (for example) optical modules, which are only available with higher interface voltages.

Another consideration in the design of the power system of the MP7 board is the fact that the board uses the AMC form-factor, and so receives power at 12V. This means that many of the voltages required on the board must be achieved by using two steps, a "bulk" step-down and a "fine" step-down (see figure 3). The physical distribution of the power-supplies on the board is shown in figure 4.

2 Multi-Gigabit Transceivers

The MP7's data-interface consists of 72 links operating in excess of 10Gbps, making full use of all the Multi-Gigabit Transceivers (MGTs) available on theVirtex-7. The MGTs on Xilinx 7-series FPGAs require three different supply voltages; 1.0V, 1.2V and 1.8V. There are, however, many

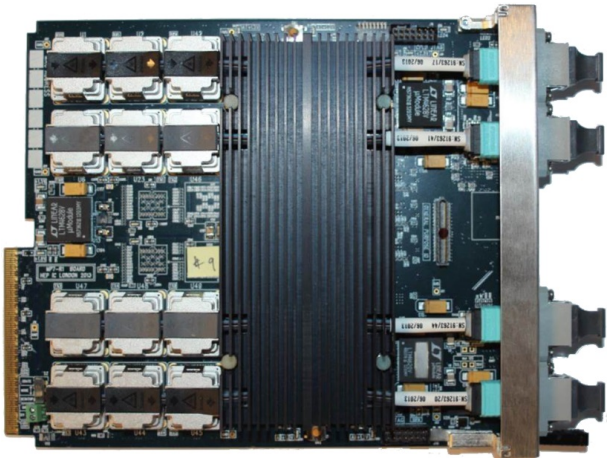

Figure 1. The Imperial MP7 processor. The MP7 is based on the Xilinx Virtex-7 FPGA (under the custom-designed heatsink) and Avago MiniPOD optical modules. The MP7 provides 72+72 optical serial links capable of operating at over 10Gbps giving a total bandwidth of 0.75+0.75Tbps.

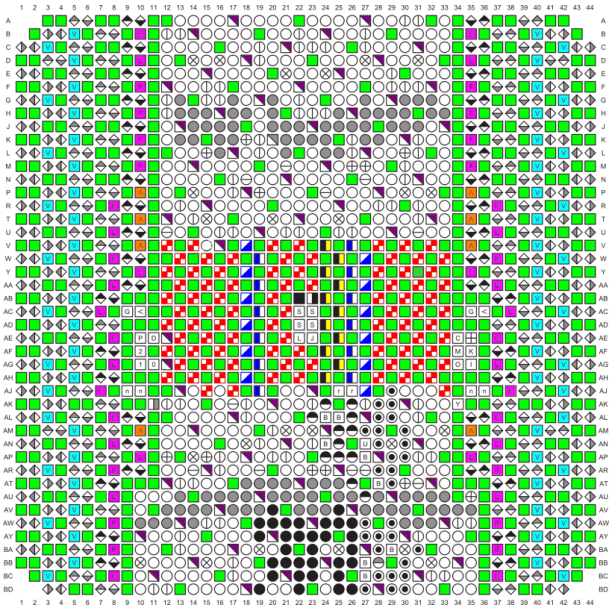

Figure 2. Pin-out of the XC7VX690T FPGA in a FFG1927 pin package. Pins coloured green are ground pins, those in other colours are power-supply pins and those in grey-scale are I/O pins: approximately half the pins are dedicated to powering the chip.

additional factors which must be taken into consideration. The noise limits and tolerances on the MGT power supplies are very tightly specified. There are also constraints on the voltage drop across the chip, such that, taking ohmic losses in the planes into account, it became apparent that the regulators needed to be as close as possible to the MGTs. Since there are two banks of MGTs in the Virtex-7, each bank needed powering independently.

Furthermore, the break-out pattern of the high-speed links, coupled with the power-supply requirements, required that the 10Gbps traces pass directly under the switch-mode regulators, and,

2013 JINST 8 C12037

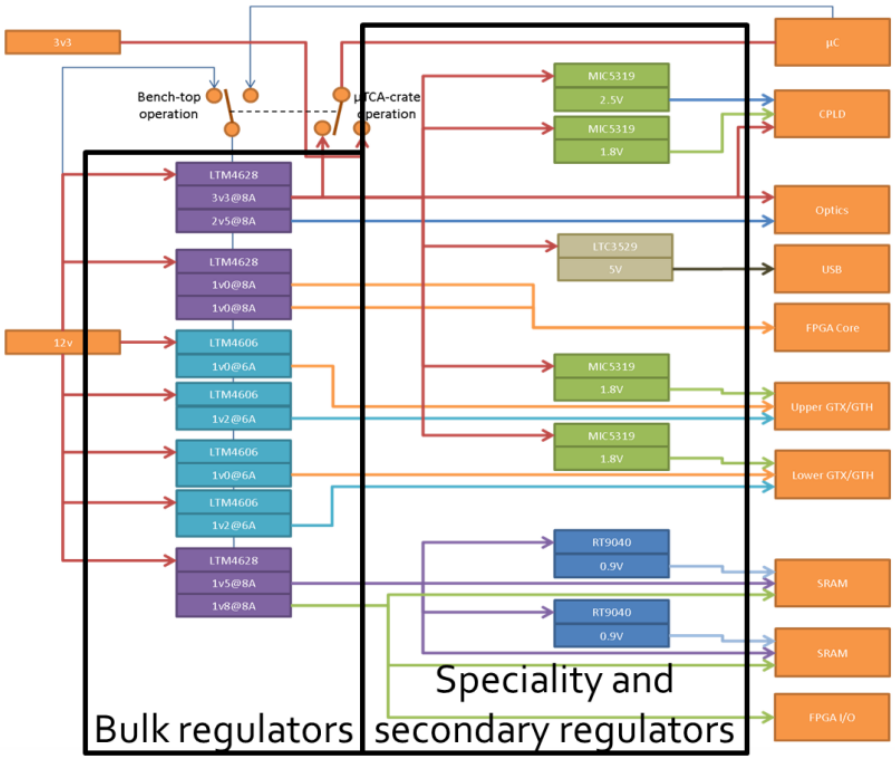

Figure 3. Power architecture used on the initial design of the MP7 processor card. A two-step architecture is used: a bulk step-down and a fine step-down.

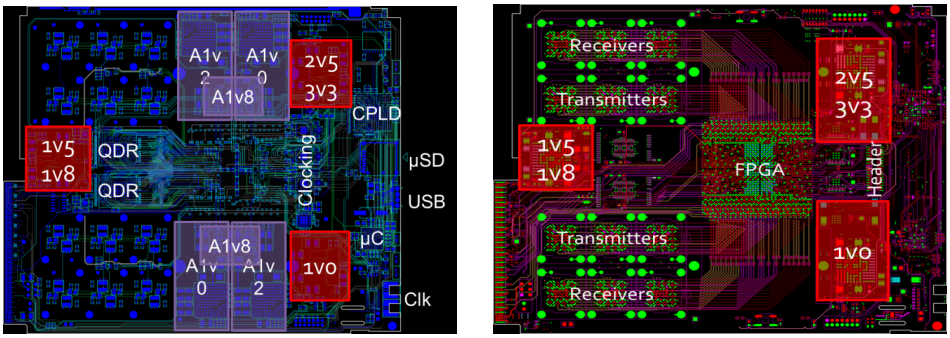

Figure 4. Distribution of power supplies on the MP7 processor board: left shows the bottom surface of the board and right shows the top surface. The red boxes indicate the bulk supplies based on the Linear LTM4620 [4] and LTM4628 [5] regulators and the purple boxes indicate the "analogue" power supplies for the multi-gigabit transceivers.

as such, the coupling of switching noise from the regulator into the high-speed traces was a concern. Our favoured component for powering the transceivers was the LTM4606 6A ultralow-EMI switch-mode regulator by Linear [6] because; it was designed for transceiver applications; it is sufficiently low profile to be fitted on the bottom of the board; and it had been successfully used in previous applications. To test the noise effect of the regulator, a commercially available kapton

2013 JINST 8 C12037

cable, consisting of high speed differential pairs over a ground plane, was modified so that one end was terminated while the other was fed into an oscilloscope. When the cable was brought into close proximity of the regulator with the differential pairs facing the regulator, noise was clearly present. The noise could be substantially reduced by simply turning over the cable so that the ground plane of the cable faced the regulator. Placing a 65μm copper tape between the cable and the regulator removed the noise entirely.

The high-speed traces were, therefore, kept in the top-half of the PCB stack-up, while the power supplies were kept in the bottom half with isolation provided by several power and ground planes, most of which were $\sim$18μm thick.

As stated previously, the MP7 was designed prior to Engineering Silicon being available and, so, the design of the MGT power supplies was based on the Xilinx Power Estimator (XPE) tool. Based on the estimation, the LTM4606 regulator should have provided $\sim$35% headroom, however, when the card was assembled, it was found that the power consumptions of the MGTs was between 30% and 220% higher than the predicted value, depending on the settings used. Understanding these effects required considerable discussion with Xilinx engineers. The MGT power estimates from subsequent XPE releases were revised upwards and now match more closely those measured.

On the initial version of the board, the LTM4606 regulators performed to, or even above, our expectation, even when operating at $\sim$25% above their stated continuous operating limit. For long-term support, however, it was felt that operating the regulators above their stated rating was undesirable, and so alternatives were sought. Naturally, the possibility of power supply changes introducing noise was a major concern, so several test cards (see figure 5a) which could be mounted on a revision-0 card were produced to test alternative power supply designs. Noise was measured both electrically and by its effect on the error-rate of the 10Gbps optical links (see figure 5b). The noise performance of all the tested designs was similar, and the design eventually chosen to replace the LTM4606 regulators was that based on the Linear LTM4601 switch-mode regulator [7]. The LTM4601 is the same size as the LTM4606 and uses similar external components, making LTM4601 the simplest replacement for the LTM4606.

3 Core power

Very large Ball-Grid Arrays (BGAs), such as the 1927-pin package used on the MP7, must necessarily puncture the board with vast forests of vias to break out their signals. The puncturing of the planes, however, increase the effective resistivity of the power planes. On the Virtex-7, the core-power pins are at the centre of the BGA, meaning that the power-planes are punctured by peripheral pins used for I/O and I/O power. The current consumed by the core-logic is not inconsiderable, so three precautions were taken; firstly, fills were placed in three separate layers to get power into the core; secondly, 1oz copper layers were preferred over $^1/_2$oz copper layers; finally, the remote sense feature of the regulator is used to monitor the observed voltage at the heart of the BGA, even though the distance is only a few centimetres.

In total (that is, core plus MGTs) the Virtex-7 FPGA on the MP7 consumes in the region of 30–40W, depending on firmware and configuration.

2013 JINST 8 C12037

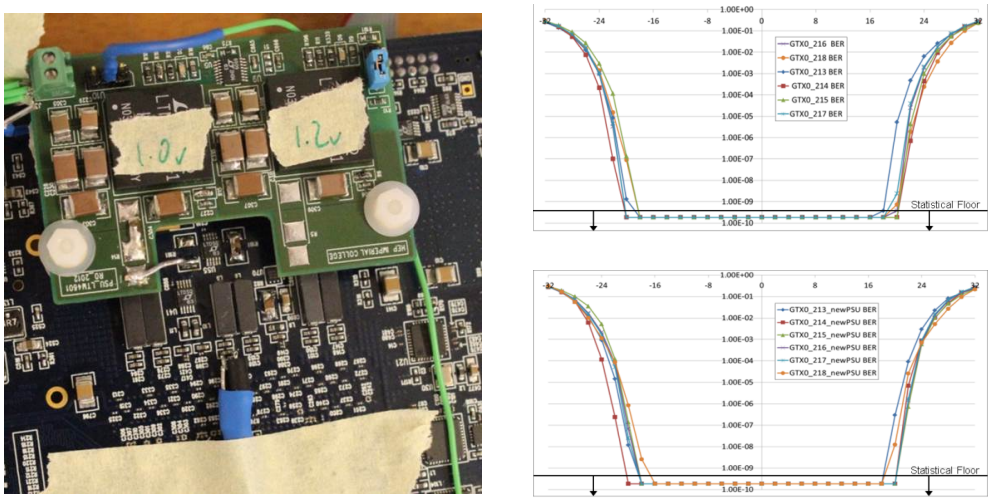

Figure 5. Left, one of the test-boards for evaluating the performance of alternative transceiver power supplies. The board shown here is the LTM4601 design chosen for the final design. The plots on the right show the bath-tub curves for a selection of the 10Gbps links using the LTM4606 (top) and LTM4601 (bottom) designs. The bit-error rate of both designs hits the statistical floor at the centre of the eye.

4 Power supply monitoring

Again, because the MP7 was designed before 7-series FPGAs were available and because of concerns about power consumption, extensive monitoring capability was included on all power supplies. For this, LTC2990 I2C four-channel voltage/current/power and temperature monitor by Linear were used, which provide sub-millivolt resolution, 1% current resolution and 1% temperature resolution, and come in a small, 10-Lead MSOP package, which was very important in the space-constrained μTCA environment. Measurements are made on both incoming supplies, on all the bulk supplies and on a subset of the secondary and speciality supplies. Furthermore, the sensors are distributed around the board allowing an approximate temperature profile of the board.

5 Heat dissipation

With an early revision of the MP7 card (see figure 6a) the 48-link design hit thermal cut-out when operated on the bench with no fans, and reached 60 $^\circ$C when operated in a μTCA crate with fans on full-power. This was a great concern and it was decided to investigate alternative heatsinks. A custom heatsink was designed (using old-school, educated guesswork and no kind of simulation) and was prototyped in-house. This custom design dissipated 40% more power than the off-the-shelf part and the FPGA temperature did not exceed 45 $^\circ$C in 48-link design. The heatsinks for the production version of the board, shown in figure 6b, were manufactured externally and anodised to provide electrical insulation and to increase the radiative heat-transfer by up-to a further 25%.

6 Conclusions

Despite the MP7 using many power saving techniques in the Virtex-7 firmware, the FPGA still requires 30–40W, which is a significant amount of power to supply and heat to dissipate. Such

2013 JINST 8 C12037

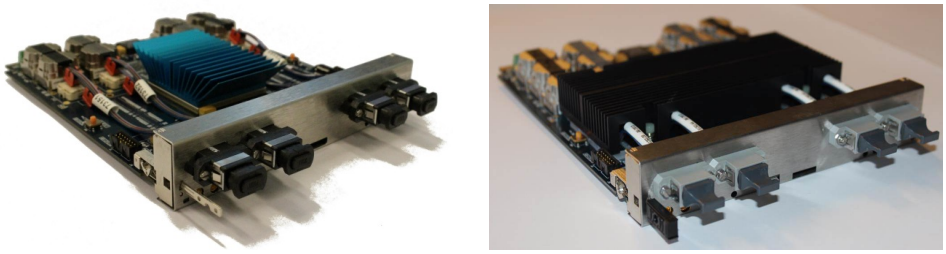

Figure 6. Left, the initial prototype MP7 used an off-the-shelf heatsink and, right, the custom-designed heatsink used for final production.

power demands also results in very large currents flowing in the card, particularly for the core supply. Designing cards to handle this power in a small form factor remains challenging.

If the total FPGA power continues to increase, cooling will become an increasingly significant problem and techniques developed for CPU cooling may have to be adopted or adapted.

Acknowledgments

Imperial College would like to acknowledge the Science and Technology Facilities Council (STFC) for funding the U.K. contribution to this research. Also, we would like to thanks Sarah Greenwood for her exceptional layout work on the MP7 processor card.

References

[1] A. Rose et al. *The MP7 and CTP-6: multi-hundred Gbps processing boards for calorimeter trigger upgrades at CMS*, 2012 *JINST* **7** C12024.

[2] *MP7 homepage*, http://www.hep.ph.ic.ac.uk/mp7, Imperial College London.

[3] Xilinx Incorporated, *7 Series FPGAs overview advance product specification*, http://www.xilinx.compsupport/documentation/data_sheets/ds180_7Series_Overview.pdf (May 2012).

[4] Linear Technology Corp., *LTM4620 dual 13A or single 26A DC/DC µModule regulator*, http://cds.linear.com/docs/en/datasheet/4620fa.pdf (Oct. 2013).

[5] Linear Technology Corp., *LTM4628 dual 8A or single 16A DC/DC µModule regulator*, http://cds.linear.com/docs/en/datasheet/4628fd.pdf (Jul. 2011).

[6] Linear Technology Corp., *LTM4606 ultralow EMI 28VIN, 6A DC/DC µModule regulator*, http://cds.linear.com/docs/en/datasheet/4606fb.pdf (Mar. 2011).

[7] Linear Technology Corp., *LTM4601/LTM4601-1 12A µModule regulators with PLL, output tracking and margining*, http://cds.linear.com/docs/en/datasheet/4601fc.pdf (Mar. 2012).

[8] Linear Technology Corp., *LTC2990 quad I2C voltage, current and temperature monitor*, http://cds.linear.com/docs/en/datasheet/2990fc.pdf (Dec. 2011).

2013 JINST 8 C12037