Program Digest, 2019 IEEE International Ultrasonics Symposium (IUS)
Glasgow, Scotland, October 6-9, 2019

A SoC-based fully configurable point-of-care ultrasound system for research purposes

Carlos J. Martín-Arguedas
Department of Electronics
University of Alcalá
Alcalá de Henares, Spain
cj.martin@uah.es

***Abstract*—This work presents a point-of-care ultrasound research system. It is composed by an analog front-end with 16 arbitrary waveform emitters and 16 low noise receivers that can be connected to arrays with up to 128 transducers. The system is controlled with a commercial SoC board, based on ZYNQ device, where the beamforming is implemented. Data are transferred via WiFi to a mobile platform for displaying.**

Keywords—Point-of-care, system-on-chip, low noise, analog-front-end, 2R-SAFT

I. INTRODUCTION

Advances in ultrasound usually require to explore, test and evaluate new acquisition algorithms, advanced digital signal processing techniques and different emission and reception strategies. This is especially relevant when researchers work on new pulse codification schemes, synthetic aperture techniques or advanced data compression approaches for point-of-care applications, where there are severe constrains on space and power consumption. Sometimes it is difficult to find suitable hardware platforms to test these techniques, and the use of high-end ultrasound machines usually doesn't solve this issue because hardware and software are protected and inaccessible to researchers. Because of this, still remains the need for open and cost-effective ultrasound systems that will allow researcher to take control of any emission and reception parameter and select whatever configuration they would need to implement their proposals [1].

Recent advances in microelectronic, and particularly the appearance of new System-on-Chip (SoC), make this goal easier [2, 3]. These devices merge in the same die the software capabilities of a microprocessor with the hardware capabilities of a FPGA, providing a new level of performance, flexibility and scalability, as well as a significant reduction in power consumption and manufacturing cost. Their heterogeneous nature allow engineers to optimize the computing architecture, extracting the maximum performance from the device and attaining a level of electronic integration never seen so far [4]. For these reasons, these component are well positioned to become the core of a new generation of high performance point-of-care ultrasound systems.

Here we present a fully configurable point-of-care ultrasound system specially designed for research purposes, and opened to the scientific community under a collaborative approach. It includes both emission and reception stages, as well as all additional elements needed to make the system completely operative. The emission stage is based on 16 high-speed sources drivers abled to emit arbitrary pulses of any length, whereas the reception stage is composed by 16 ultra low noise channels sampling up to 80MHz and 12 bits per sample. A high-voltage multiplexing system allow to connect this front-end to 128 pins connector, where commercial probes can be plugged in. All power rails are included on board, so only a 12v power supply is required to make it run. The acquisition system is controlled from a low-cost ZC702 Xilinx demo board [5], which integrates a ZYNQ device where user firmware can be easily implemented. In this particular setup the ZYNQ includes a back-end system that manage the emitters and receivers, generate clocking signals, process data flows coming from analog-to-digital converters and compose the final images, which are transferred to a smartphone thru a WiFi connection to be displayed in real-time.

II. SYSTEM DESCRIPTION

A. Front-end board

The analog front-end includes the emission and reception modules, and additional components to power them. The emission stage is based on 16 high-speed sources drivers (MD2134, Supertex Inc.) capable to emit arbitrary waveforms of any length. Every driver is connected to a central-tap RF transformer, which is powered with a 90 volts external power supply to avoid switching noise contamination. Thus, +/-90v amplitude waveforms can be easily obtained in the secondary output.

The reception stage is composed by two low noise 8-channels analog-front-ends (AFE5807, Texas Instruments), sampling up to 80MHz and 12bits. These parts include low noise amplifier, voltage controlled attenuator, programmable gain amplifier, low pass filter, and high-speed double data rate LVDS serializes, so a maximum throughput of 15Gbps can be generated. Ultra low noise clock buffers are included on board to generate and distribute clock signals to the receivers. A set of four 32 inputs-16 outputs high voltage analog switches (HV2801, Supertex Inc.) allow connecting the front end to medical probes with up to 128 transducers. The multiplexing systems can be easily configured using a SPI port, so different activation schemes can be easily implemented. These multiplexers are powered up with a non-switching +/-100v power supply included on board. To protect the reception stage from EMI noise, isolated buck converters and ultra-low noise

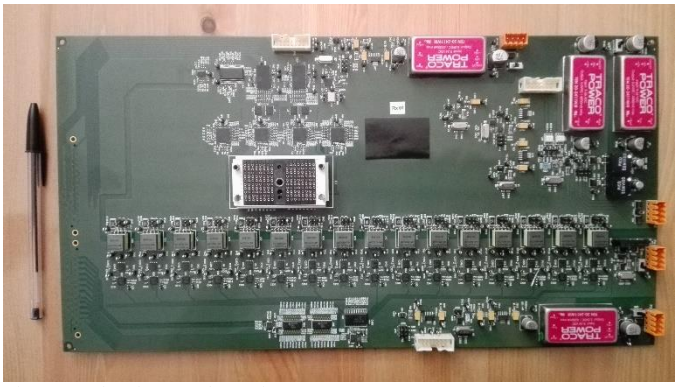

Fig. 1. Analog front-end top view

LDOs have been included in all rails. A set of sensors to get accurate measures of voltage, current and power consumption in every rail are also included. Data acquired with these sensors can be read and transferred to a PC in real-time using a Linduino board [6].

Fig. 1 shows an image of the platform, where the Hypertac connector to link with commercial probes can be seen in the center of the board.

B. *Back-end board*

In order to processing raw signals coming out from the analog front-end, a ZC702 Xilinx demo board is used. It includes a ZYNQ 7020 device. This is a cost-optimized SoC with an ARM Cortex-A9 Dual Core processing system (PS) and an Artix-7 programmable logic (PL), all together in a single device [7]. The ZC702 board also includes 1GB of DDR3 memory, plus additional interfaces to connect with external peripherals (HDMI, CAN, I2C, Ethernet, etc). It also has an advanced powering system, which can be monitored using a PMBus, in order to extract information about consumption.

Front-end and back-end boards connect each other by means of two high density FMC-LPC connectors located at the bottom of the first one. Fig 2. Shows an image of the entire system.

In this particular case, the ZC702 board has been enough to implement the entire system. However, if additional capabilities are demanded, the ZCU106 demo board, which includes the high-range ZU7EV Zynq UltraScale+ MPSoC, can be also connected to the front-end, increasing considerably the system capabilities.

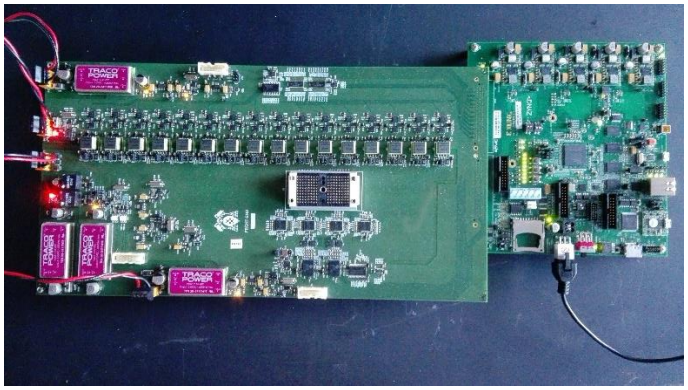

Fig. 2. Front-end and back-end boards connected

III. Processing chain implementation

In this particular case, a 2R-SAFT technique has been used to conform the images [8]. This is a very straightforward synthetic aperture technique that allows to remove grating lobes in ultrasound images and simplify the acquisition process. This is particularly useful in point-of-care applications, where a tradeoff between resources and performance is required to minimize hardware size. Here, the entire beamforming process is performed within the ZYNQ device, avoiding the need to use an external GPU. To do that, implementation takes advantage of the NEON coprocessors present in the ARM cores. There are two in total, one for each core, and they are based on a Single Instruction Multiple Data architecture, similar to GPUs [9]. Typically these coprocessor reach an 8x speedup.

The processing blocks are divided in two sections. Thus, deserialization, codec decorrelation and DC filtering are performed on PL, whereas I-Q calculus, beamforming and logarithmic compression are performed on PS. Controllers to write and read AFEs registers, configure multiplexors and send stored waveforms to the pulsers are also implemented in PL.

In order to extract the maximum performance from the SoC, the Cortex-A9 processor works in asymmetric multiprocessing configuration (AMP), where both cores works independently and run their own operative system or bare-metal application. Here, a PetaLinux runs on CPU0, which is the master and responsible for system initialization, startup the CPU1 and interact with the user. Simultaneously, a bare-metal application runs on CPU1, which is responsible to control the signal processing blocks and beamforming images. The inter-processor communication is implemented using OpenAMP [10], a Xilinx library that provides the software components to create a message communication between cores using shared memory and inter-cores interruptions.

Basically, data coming from the analog front-end enter directly into the PL. Here, they are processed and filtered, and then transferred to the external DDR memory using a DMA port. Then, CPU1 read data from memory and perform the beamforming, copying back the resultant image to the DDR. Then, CPU0 read image from memory and transfer to the smartphone using a WiFi link, which is implemented over the USB3.0 peripheral included in the PS.

IV. SoC Resources and Performance

To evaluate the system performance, a 64 elements, 3.5MHz phased array (IMASONIC, France) was used to acquire images with the 2R-SAFT technique. A total of 127 signals with 8196 samples per signal (~8cm depth in water) were obtained to compose every image. Chirp signals with ~40% relative bandwidth were used in emission. In reception, dynamic focusing was applied to conform images with 512x512 pixels that were transferred to a smartphone for displaying.

Table I summarize logic utilization after place and route optimization.

TABLE I. DEVICE UTILIZATION SUMMARY

Logic	Used	Available	Utilization
Slice registers	14.745	106.400	13'86%
Slice LUTs	10.832	53.200	20'36%
Block RAM Tiles	76'5	140	54'64%
DSPs	0	220	0%
IDELAYCTRL	1	4	25%
IDELAY2	16	200	8%
ISERDES	32	200	16%

More demanded resources are Block RAM Tiles and ISERDES. The first are used for FIFO implementation and storing filters coefficients mostly, so this value is strong dependent with the number of samples acquired in every fire and the length of the codes in emission. On the contrary, ISERDES are required to deserialize the LVDS input channels. The AFE5807 has LVDS DDR interfaces, so there are different bit data on the positive as well as negative edge of the bit clock. For this reason, two ISERDES are needed to deserialize every input channel. Definitely the number of ISERDES is limited in ZYNQ architectures, so it can be a drawback for increasing the total of input channels far away from 96. However, this is a number quite reasonable for a point-of-care platform. This increment can be easily achieved whit minimum changes in the front-end. Other logic resources, like registers, LUTs or DSPs are not a problem.

Power dissipation of the entire system (analog front-end + SoC board) is ~3'72W during typical operation, where 0'92W correspond to the SoC board, and around 2'8W to the front-end. There, AFEs demands the majority of the power, requiring up to 2'2W. High voltage pulsers consume 0'4W. With this implementation a frame rate of 30 images per second is attained.

V. CONCLUSIONS

New advances in microelectronic, like System-on-Chips or advanced feature converters, make possible the design of high performance ultrasound systems for point-of-care applications. Here, a fully configurable ultrasound system aimed at the research community to reach that goal has been presented. It can be used to test different system architectures and implementations, and can be configured in many different ways. One particular implementation has been presented. Results show how the entire system can be easily implemented in just one SoC device.

REFERENCES

[1] E. Boni *et al.*, "Ultrasound open plarforms for next-generation imaging technique development," IEEE Trans. Ultrasonic, Ferroelectric and Frequency Control, vol. 65, no. 7, pp. 1078-1092, 2018

[2] J. Kang *et al*., "A System-on-Chip Solution for Point-of-Care Ultrasound Imaging Systems: Architecture and ASIC Implementation," IEEE Transactions on Biomedical Circuits and Systems, vol. 10, no. 2, pp. 412-423, April 2016.

[3] P. Govindan *et al.*, "Reconfigurable and programmable System-On-Chip hardware platform for real-time ultrasonic testing applications," *2015 IEEE International Ultrasonics Symposium (IUS)*, Taipei, 2015, pp. 1-4.

[4] L. H. Crockett *et al.*, The Zynq book: Embedded processing with the ARM Cortex-A9 on the Xilinx Zynq-7000 all programmable SoC, Strathclyde Academic Media, UK, 2014

[5] https://www.xilinx.com/products/boards-and-kits/ek-z7-zc702-g.html

[6] https://www.analog.com/en/design-center/evaluation-hardware-and-software/linduino.html

[7] https://www.xilinx.com/support/documentation/data_sheets/ds190-Zynq-7000-Overview.pdf

[8] C. J. Martin-Arguedas *et al.*, "An ultrasonic imaging system based on a new SAFT approach and a GPU beamformer," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 59, no. 7, pp. 1402-1412, July 2012.

[9] https://developer.arm.com/architectures/instruction-sets/simd-isas/neon

[10] https://www.xilinx.com/support/documentation/sw_manuals/xilinx2018_2/ug1186-zynq-openamp-gsg.pdf