

PC-Based Modular Digital Ultrasound Imaging System

Amr M. Hendy, Mawia Hassan, Rania Eldeeb, Dina Kholy, Abou-Bakr Youssef and Yasser M. Kadah
Biomedical Engineering Department, Cairo University and IBE Tech, Giza, Egypt
E-mail: ahendy@ibetech.com

Abstract—With the availability of high-end integrated analog front-ends, distinction between different digital ultrasound systems is determined almost exclusively by their software component. Moreover, PC-based implementations for sophisticated medical imaging technologies have emerged where powerful multi-core computational ability replaces expensive embedded systems. The objective of this work is to develop a modular low-cost PC-based digital ultrasound imaging system that has almost all of its processing steps done on the PC side.

Keywords—PC-based ultrasound; digital beamforming; embedded signal processing; PCI-Express bus.

I. INTRODUCTION

With the growing availability of high-end integrated analog front-end circuits, distinction between different digital ultrasound imaging systems is determined almost exclusively by their software component. Efficient implementations of digital ultrasound systems rely on embedded digital signal processing on FPGA with data conversion from oversampled 1-bit delta-sigma analog-to-digital converters (ADC) to minimize the number of lines going into the FPGA. However, using LVDS interface protocol allows a serial output with drivers on FPGA to recover the parallel data. This alleviates the need for designing the sampling and the signal recovery filters while maintaining an optimal performance at significantly lower power consumption. Moreover, PC-based implementations for sophisticated medical imaging technologies have emerged where powerful multi-core computational ability replaces expensive embedded processing systems. This approach reduces the overall cost of the system as well as the speed of system development. The objective of this work is to develop a modular low-cost PC-based digital ultrasound imaging system that has almost all of its processing done on the PC side.

The main contribution of this work is to move almost all know-how related software components into the PC side of the system. The system block diagram for the system is shown in Fig. 1 where the data are collected and interfaced to the PC via a PCI-express bus through a Virtex-5 FPGA (Xilinx, Inc.). The use of several of these modules is possible through the use of multiple lanes of this interface bus.

II. PREVIOUS METHODS

Several papers addressed the issues involved in digital beamformer design [1-2] including the description of its main components. Embedded digital beamforming was initially

done using Application-Specific Integrated Circuits (ASICs) [1]. Among the issues addressed is the sampling issue where a number of oversampled implementations were proposed to allow for artifact free conversion of ultrasound signals to the digital domain [1]. The phase aberration effect has also been described and its effects on ultrasound beamforming have also been assessed [1]. Many articles also addressed the digital signal processing algorithms that can be used in digital beamforming signal demodulation [2]. The issues involved in real-time digital ultrasound imaging are described in [3] including estimates for the computational complexity involved in each part of the ultrasound digital signal processing chain.

An interesting approach for efficient digital beamforming was presented in [4]. In their technique, a compact medical ultrasound beamformer architecture that uses oversampled 1-bit analog-to-digital converters (ADC) is presented. Sparse sample processing is used, as the echo signal for the image lines is reconstructed in 512 equidistant focal points along the line through its in-phase and quadrature components. That information is sufficient for presenting a B-mode image and creating a color flow map. The high sampling rate provides the necessary delay resolution for the focusing. The low channel data width (1-bit) makes it possible to construct compact beamformer logic. The signal reconstruction is done using finite impulse response (FIR) filters, applied on selected bit sequences of the delta-sigma modulator output stream. The approach allows for a multi-channel beamformer fitting in a single field programmable gate array (FPGA) device. A 32-channel beamformer is estimated to occupy 50% of the available logic resources in a commercially available midrange FPGA, and to be able to operate at 129 MHz. Simulation of the architecture at 140 MHz provides images with a dynamic range approaching 60 dB for an excitation frequency of 3 MHz. In spite of the interesting approach used, a number of significant issues can be raised in this method. First, the verification of the proposed methodology was done using computer simulations. The actual implementation of this method would involve dealing with such tricky issues as interfacing hardware to FPGA, printed circuit board design to minimize noise and cross talk between tracks of high-speed digital signals, as well as maintaining the power supply quality for the analog front-end with many fast switching digital signals in close proximity. Also, the justification of using oversampled 1-bit analog-to-digital converters in this work was to minimize the number of lines going into the FPGA. However, with the current analog-to-digital converter and FPGA technologies, another alternative has been the focus of

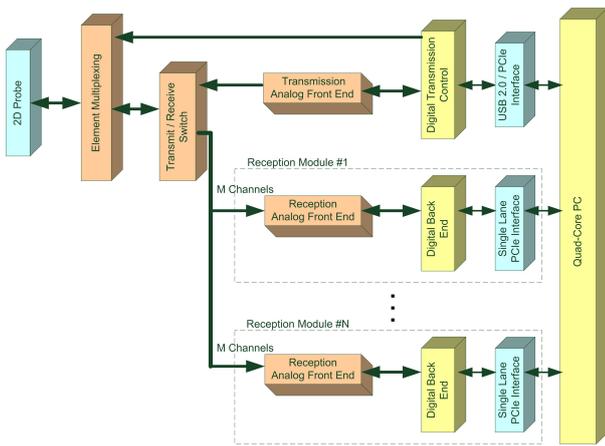


Fig. 1: Block diagram of proposed digital ultrasound imaging system

most fast analog acquisition implementations in the past few years. This alternative is the use of serial low voltage differential signaling (LVDS) interface protocol, which allows the output from an ADC to come as a serial bit stream with drivers on both the ADC and the FPGA to recover the parallel data. This alleviates the need for designing the sigma-delta sampling and the signal recovery filters for the oversampled 1-bit data stream while maintaining an optimal performance at significantly lower power consumption.

In this work, we take the following comments into account in planning for the new digital beamformer design. Moreover, we offer two additional advantages over this previous work, namely, the new integrated analog front-end components and the more powerful FPGA series that became available after this previous work was done.

III. METHODOLOGY

If we consider the process of collecting digital samples from a line with the maximum scanning depth of 22 cm, the number of samples collected by the 32 channel digital beamformer at the desired sampling rate of 50 MSPS at 12 bit resolution is about 894 kBytes. For a real-time acquisition of lines, the data rate entering the digital beamformer may reach up to 3.2 GBytes. It is not a trivial task to handle such huge amounts of data. Therefore, it is desired to somehow reduce such data without compromising the performance of the system. The basic idea we developed to do that is based on the following points. First, the actual bandwidth of B-mode ultrasound imaging data in its quadrature form does not exceed a few MHz even with high frequency probes. Moreover, for Doppler ultrasound data, such bandwidth may even be as lower. The sampling done in hardware produces real samples and not the quadrature components representing the analytical form of each sample. It is difficult to have analog quadrature demodulation in the planned design because

of the integration of the analog front end up until the ADC within one chip. Also, ultrasound signals before this chip are too weak to process with this scheme. This of course in addition to the limitation of analog quadrature demodulation assuming narrowband nature for signals. This means that the derivation of the analytical signal in our design has to be on the digital side. Note that the Hilbert transform can be used to derive the analytical signal from its real part. With the analytical form of the signal, the signal can be downsampled to its bandwidth without aliasing while keeping the phase information intact for further phase-sensitive processing as wideband beamforming or Doppler shift estimation. Hence, for the ultrasound imaging situation, one can downsample a Hilbert transformed single channel signal from 50 MSPS to a mere 5 MSPS without losing any information. A number of authors suggested the use of digital finite impulse response (FIR) filter approximations to implement the Hilbert transformation. Our initial experiments on this with Matlab resulted in simple filters that can be used to do that with accuracy/filter length trade-off. Nevertheless, Even for a 20-tap FIR filter implementation of the Hilbert transformation, it is still a cumbersome task to implement several filters with an input and output data rates of 50 MSPS. If we plan to downsample the analytical form of the data after the filter, it makes no sense to compute 50 MSPS in the output when we indeed plan to throw away nearly 95% of them. In particular, the FIR filters are designed in such a way to do both the filtration and downsampling at the same time and we need to compute only 5 MSPS per channel rather than 50 MSPS. A block diagram of the system is shown in Fig. 1.

IV. PC INTERFACING

The current PC technology offers a number of very fast external interfaces that are intended to support fancy 3D gaming requirements [5]. The fastest of these interfaces now is the Peripheral Component Interconnect Express bus (PCI-Express or simply PCIe) interface that is utilized now in mainstream PCs to interface graphics cards. This interface is a very fast serial interface that is capable to deliver a transfer rate of 250 MB/s (Gen 1.x), 500 MB/s (Gen 2.0), or 1 GB/s (planned for Gen 3.0) per single lane. It allows also as many lanes as 16 to be used, delivering a very powerful of up to 16 GB/s data transfer capacity for a single device. With multiple PCIe slots, multiple devices working at this transfer rate can be used in tandem. Hence, the data transfer rate for ultrasound systems is now within reach using this technology. At the time when this work started, PCIe Gen 1 was available and goal was to develop an 8 lane PCIe Gen 1 interface based on a Xilinx Virtex 5 FPGA device that was planned to deliver nearly 2 GB/s transfer throughput. This is sufficient for a 16 channel reception module.

In using PCIe, given its serial interface protocol, several issues must be taken into consideration in designing the transfer. Examples include maximum payload size and variation of packet efficiency with payload size that may reduce the transfer bandwidth by nearly 20%. A major issue arises also in that the data transferred has to be written to the computer memory further processing, which adds a bottleneck for the PC software part of the system. Therefore, we decided to address the problem of lossless reduction of ultrasound data bandwidth.

V. DATA PROCESSING STRATEGY

The goal of this section is to describe the method used to reduce the required raw ultrasound data transfer bandwidth while maintaining the phase information. The processing steps are shown in Fig. 2 for single channel data (with other channels using the same architecture). Assuming that the sampling rate of each channel is N Sa/s, the acquired samples represent an oversampled version of the real part of the signal. Given the frequency characteristics of most ultrasound probes having around 60% bandwidth around the center frequency, the frequency spectrum of the signal is sparse with a significant part of the spectrum having negligible signal components. So, it is possible to exploit this signal characteristic to make a dramatic reduction in signal bandwidth while maintaining the original information intact. Using a discrete FIR Hilbert transform filter, the analytic signal can be computed with only the positive side of the original signal spectrum. Hence, such analytic signal can be downsampled to only the bandwidth of the analytic signal, which is much smaller than that of the real signal.

Note that in the Hilbert transformation exact implementation, it acts like an ideal filter that removes all the negative frequencies and leaves all positive frequencies untouched. In its Matlab implementation, it is done using frequency domain filtering after the discrete Fourier transformation (DFT) is applied to the real signal. This ideal implementation affords high performance for broadband signals but at the same time requires high computational complexity (or FPGA resources in embedded implementations). Therefore, a number of authors suggested the use of digital finite impulse response (FIR) filter approximations to implement the Hilbert transformation. Even though simple filters can be used with accuracy/filter length trade-off, the computational complexity of working on the original high sampling rate is still not practical for direct implementation. For example, for a 20-tap FIR filter implementation of the Hilbert transformation, it is still a cumbersome task to implement several filters with an input and output data rate of 50 MSa/S. This complexity can be significantly reduced by combining the downsampling and filter implementation together. If we plan to downsample the analytical form of the data after the filter, it makes no sense to compute 50 MSPS in the output when we indeed plan to throw away most of them. Therefore, the Hilbert FIR filters are designed as multirate filters in such a way to do both the filtration and downsampling at the same time. For the above

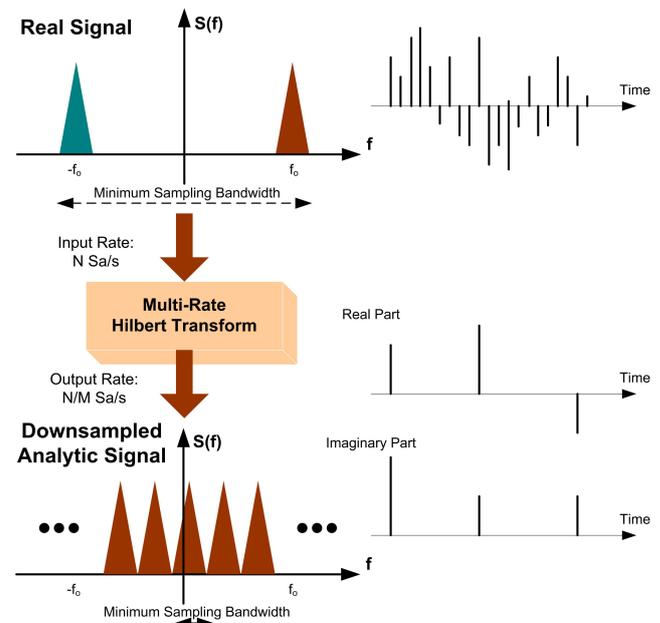


Fig. 2: Sampling strategy

example, for an analytic signal bandwidth of 5 MHz, the filter outputs only 5 MSa/s per channel rather than 50 MSa/s.

The analytic signal consists of the original real signal decimated by the factor of choice and the output of the multirate Hilbert filter that outputs a decimated version of the imaginary part of the signal. The Hilbert filter is implemented in the common optimized form whereby the zero tap coefficients are not computed and therefore an order L filter uses only $L/2$ multiplications. Also, the coefficients are implemented as 16-bit signed integers. Given that the same coefficients are used for all channels, it is possible to use a longer FIR Hilbert filter for multiple channels in an interleaved manner. In this case, the signal outputs for different channels sustain small delay differences that can be easily compensated in the further beamforming stage of the system.

VI. PC PROCESSING

Given that the number and geometry of ultrasound scan lines vary with probe and in all cases are different from what practical image display requires. It is necessary to perform a scan conversion step to reconstruct the ultrasound image. In its most general form, this step takes raw ultrasound lines along with their geometrical properties (i.e., direction, spacing, etc.) and estimates a rectilinear array of values representing the image that is compatible with the display device. The specifications of our design require the image to be a 512×512 matrix of values. This step is done in real-time on the PC used to display the image.

VII. EXPERIMENTAL VERIFICATION

The new system was implemented using Analog Devices ultrasound analog front end chipset interfaced to a Virtex-5 FPGA with an 8-lane PCIe interface. The system was plugged into a Pentium 3.0 GHz Core 2 Quad PC with 8 GB memory. The system was used to acquire data from a resolution

phantom. The sampling rate was 50 MHz and the number of channels used acquired was 32. The scan depth was 6 cm. The PCIe interface bandwidth did not allow real-time raw data transfer to the PC for processing because of several efficiency problems that did not permit the system to reach its peak transfer rate (such as packet size). The transfer rate was significantly boosted by using the compression strategy above. In Fig. 2, the data from a single channel centered above one of the pins in the scanned phantom is shown. The results of applying the Hilbert transform to compute the analytic signal is shown in Fig. 3 where the frequency spectrum of the signal is computed and shows a single sided spectrum for the analytic signal. The spectrum after applying a divide-by-8 downsampling step is shown in Fig. 4 before and after the required demodulation. Notice that the downsampling was preceded by an antialiasing filter implemented as an FIR digital filter and combined with the Hilbert transform to reduce the noise superimposed from other frequencies in the downsampling step. To show that this processing strategy did not cause information loss, the real signal was reconstructed and compared to the original as shown in Fig. 5. The error between the two signals was lower than 1% for this particular experiment but increases to nearly 10% without the antialiasing filter. Once the data were in the PC memory, reconstruction of final image was done using spatially variant filtration of the received line data implemented using a look-up table. This allows the utilization of the parallel programming offered by the quad-core processor used in the system. Real-time reconstruction frame rates were achievable on the preliminary version of this system. In Fig. 6, a magnified 2cm \times 1.5cm part of the image containing the resolution pins is shown.

VIII. CONCLUSIONS

A technique that expands the utility of current high-performance PCs to replace the high-cost embedded processing in digital ultrasound systems was studied. The new system has the potential to lower the cost and speed up the development, thus offering new opportunities for more cost-effective systems.

REFERENCES

- [1] B.D. Steinberg, "Digital beamforming in ultrasound," *IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control*, vol. 39, no. 6, pp.716-721, 1992.
- [2] C. Fritsch, M. Parrilla, T. Sanchez, O. Martinez, "Beamforming with a reduced sampling rate," *Ultrasonics*, vol. 40, pp. 599-604, 2002.
- [3] C. Basoglu, R. Managuli, G. York, and Y. Kim, "Computing requirements of modern medical diagnostic ultrasound machines," *Parallel Computing*, vol. 24, pp. 1407-1431, 1998.
- [4] B.G. Tomov and J.A. Jensen, "Compact FPGA-Based Beamformer Using Oversampled 1-bit A/D Converters," *IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control*, vol. 52, no. 5, pp. 870-880, May 2005.
- [5] A. Goldhammer and J. Ayer, Jr., "Understanding the performance of PCI Express systems," *White Paper: Xilinx Virtex-4 and Virtex-5 FPGAs*, WP350 (v1.0) July 28, 2008.

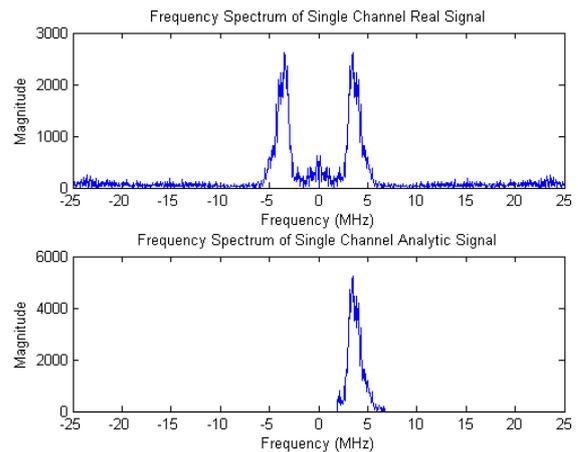


Fig. 3: Frequency spectrum of real and analytic signal

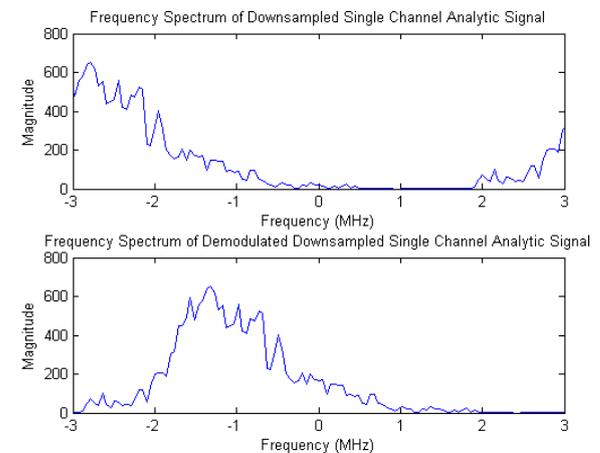


Fig. 4: Frequency spectrum of downsampled analytic signal before and after the demodulation step

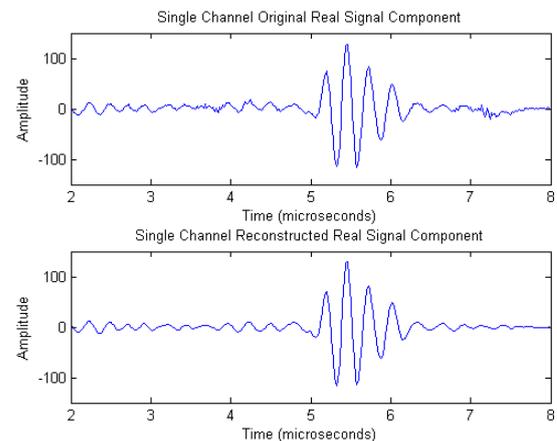


Fig. 5: Reconstructed and original real signals



Fig. 6: Zoomed part of the reconstructed image

This work was supported by Grants RDP_03_07_40 and RDP_03_07_41 from the Industrial Modernization Center, Ministry of Scientific Research and IBE Tech.