

Towards Ultrasound Everywhere: A Portable 3D Digital Back-End Capable of Zone and Compound Imaging

Aya Ibrahim ¹, Shuping Zhang ², Federico Angiolini, Marcel Ardit ³, *Member, IEEE*, Shinji Kimura, Satoshi Goto, *Life Fellow, IEEE*, Jean-Philippe Thiran, and Giovanni De Micheli ⁴, *Fellow, IEEE*

Abstract—Ultrasound imaging is a ubiquitous diagnostic technique, but does not fit the requirements of the telemedicine approach, because it relies on the real-time manipulation and image recognition skills of a trained expert, called sonographer. Sonographers are only available in hospitals and clinics, negating or at least delaying access to ultrasound scans in many locales—rural areas, developing countries—as well as in medical rescue operations. *Telesonography* would require an advanced imager that supports three-dimensional (3-D) acquisition; this would allow untrained operators to acquire broad scans and upload them remotely for diagnosis. Such advanced imagers do exist, but do not meet several other requirements for telesonography, such as being portable, inexpensive, and sufficiently low power to enable battery operation. In this work, we present our prototype of the first portable 3-D digital ultrasound back-end system. The prototype is implemented in a single midrange Xilinx field programmable gate array (FPGA), for an estimated power consumption of 5 W. The device supports up to 1024 input channels, which is state of the art and could be scaled further, and supports multiple image reconstruction modes. We evaluate the resource utilization of the FPGA and provide various quality metrics to ascertain the output image quality.

Index Terms—Field programmable gate arrays, telesonography, ultrasound imaging, ultrasound quality metrics, volumetric ultrasound reconstruction—3D imaging.

I. INTRODUCTION

ALTHOUGH medical *Ultrasound (US)* imaging has been in use since the 1950s, it is still an active area for research. Medical US is a very popular imaging method due to

Manuscript received February 14, 2018; revised April 6, 2018; accepted April 11, 2018. Date of publication June 4, 2018; date of current version October 19, 2018. This work was supported by the Swiss Confederation funding through the UltrasoundToGo project of the Nano-Tera.ch initiative. This paper was recommended by Associate Editor M. Martina. This paper was presented in part at the 13th IEEE Biomedical Circuits and Systems Conference, Turin, Italy, October, 2017. (*Corresponding author: Aya Ibrahim.*)

A. Ibrahim, F. Angiolini, M. Ardit, and G. De Micheli are with the École Polytechnique Fédérale de Lausanne (EPFL), 1015 Lausanne Switzerland (e-mail: aya.ibrahim@epfl.ch; federico.angiolini@epfl.ch; marcel.arditi@epfl.ch; giovanni.demicheli@epfl.ch).

S. Zhang, S. Kimura, and S. Goto are with Waseda University, Tokyo 169-8050, Japan (e-mail: zhangshuping@fuji.waseda.jp; shinji_kimura@waseda.jp; goto@waseda.jp).

J.-P. Thiran is with the École Polytechnique Fédérale de Lausanne (EPFL), 1015 Lausanne, Switzerland, and also with the Department of Radiology, University Hospital Center (CHUV) and University of Lausanne (UNIL), Switzerland (e-mail: jean-philippe.thiran@epfl.ch).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TBCAS.2018.2828382

its comparatively low-cost, safe, simple, and precautions-free scanning routine. In contrast, other imaging techniques like X-ray and *Computed Tomography (CT)* are based on ionizing radiation, and *Magnetic Resonance Imaging (MRI)* is more expensive, more oppressive, and requires strict precautions. US is considered the go-to imaging technique in women's health (obstetrics, gynecology), and is broadly adopted in many other different applications like cardiology, ophthalmology, urology, angiology, etc. Nonetheless, US is still unavailable whenever no trained specialist is available to acquire the scans. With current technology, a well-trained sonographer must finely manipulate the probe until the desired 2D body section is identified and captured. With such a high dependency on the presence of a sonographer to acquire scans, it is consequently very difficult to reap the benefits of telemedicine approaches in US domain, a goal that would be called *telesonography*. Telesonography is the concept of remotely diagnosing US scans by allowing the decoupling of acquisition - something that would be done by a lightly-trained nurse, a family doctor, etc. - from diagnosis. Telesonography would create the opportunity to enable US diagnosis in rural areas, in rescue situations (e.g., medical helicopters or ambulances), and in enclosed spaces like air-crafts, battle-fields, ships, etc. Unfortunately, telesonography has not been enabled yet and, until today, US is only available in hospitals and specialized cabinets.

3D US imaging, a technology originally developed and still used mainly for cardiology applications, can solve this problem by acquiring volumetric images at once. The 3D nature of the scans relaxes the precision constraints on the positioning of the probe, potentially allowing untrained or lightly-trained operators, e.g., paramedics, to acquire scans on the field. Either in real-time or subsequently, these volumes could be uploaded to a hospital, where a resident sonographer would reach a diagnosis. Unfortunately, today's 3D US systems are very expensive, stationary and power-consuming, as the result of the enormous computation requirements of volumetric US reconstruction. The result is that, once again, 3D US is still limited to hospital environments [1]–[3]. To solve this vicious cycle, it is necessary to design a 3D US imager that has good image quality but is cheaper and portable, and ideally can operate on batteries. Point-of-care devices do exist today [4]–[9] but they are limited to 2D, which again introduces the strict dependency on having a sonographer available for the scan acquisition.

A. Ultrasound Imaging Working Principle

In US imaging, an array of piezoelectric elements (*probe*) emits a *Radio-Frequency (RF)* (2–20 MHz) acoustic wave into the *Region Of Interest (ROI)*, which ideally should consist mostly of soft tissues. By using a probe comprising a one-dimensional array of elements D , a 2D body slice can be insonified; with a two-dimensional probe, a quasi-pyramidal 3D volume can be scanned. Tissue inhomogeneities scatter the incident waves, acting as secondary point sources (*scatterers*). The reflected echoes, still in the RF range, are received back by the same probe and the signals are then sent over a cable to a backend system for image reconstruction. The processing pipeline comprises different stages, that vary from a system to another. Nonetheless, the most essential and the core of any US processing pipeline is a process called *beamforming (BF)*. BF reconstructs an image by defining the reflectivity of the tissue (scatterers) at numerous locations, called focal points. This is accomplished by summing the returned echoes according to a delay profile that represents the time taken by a US wave to travel from an origin O to a focal point $S \in V$ and be backscattered to a transducer element $D \in 1, \dots, N$. The reflectivity correlates to the amplitude of the summed echoes and is ultimately rendered as a darker or brighter pixel in a greyscale image. The kernel of the BF process is the calculation of one delay profile per focal point. This includes another important step named *apodization*. Apodization is a weighting function applied to reduce echo contributions from the sidelobe region of the transducer aperture, which degrade both contrast and lateral resolution in the final image. The BF process can be expressed mathematically by (1) where $e(D, t_p)$ is the echo received by an element D at a time-of-flight t_p , and $w(t, D)$ is the weighting applied to the echoes sampled at time t (refer to Section III-B3).

$$s(S) = \sum_{D=1}^N e(D, t_p(|\vec{OS}|) + t_p(|\vec{SD}|))w(t, D), \forall S \in V, (1)$$

The beamformed image then undergoes a post-processing and visualization step to make it displayable on a screen and interpretable by the human eye. In particular, *Scan Conversion (SC)* is a necessary transform from polar to Cartesian coordinates, and *log-compression* is a brightness mapping from the original high-dynamic-range signal onto a logarithmic scale adapted to human vision.

In 3D imaging, both the BF and the SC processes represent the main bottlenecks and the most computationally demanding units. They are the main reason for the high power, cost, and area requirements of 3D US systems. As an example, in 3D imaging, a straightforward implementation of (1)-for a reasonable imager configuration and frame rate - leads to the need to calculate about 105 trillion delays/s [10]–[12]. Each delay value is calculated as the Euclidean distance from O to S then S to D , i.e., requiring two square roots. This is obviously extremely challenging to fit into a portable, inexpensive, and battery-powered device. Further, image enhancement techniques are widely used in medical practice and should be supported, at an additional complexity cost. This includes for example compound imaging [13], whereby the final image is obtained by combining multiple ROI

reconstructions (refer to Section III-E). For the same overall frame rate, this technique has a linear computation overhead, i.e., if the compounding count is 5, then 5 times as many delay calculations will be required. Other techniques, such as zone imaging, improve resolution, but also entail additional complexity.

To keep the processing complexity manageable, existing medical systems, both commercial and scientific, pre-process the echoes from the US probe, merging the signals from large numbers of transducer elements onto far fewer (few hundreds) channels. This is commonly achieved by either combining the signals of transducer elements and mapping them to fewer channels by a technique named *analog pre-BF* [14]–[17], or activating fewer transducer elements in patterns at a time using *multiplexing* [18], [19] or *sparse 2D-arrays* [20], [21]. However, the deliberate discarding of input signal information inevitably comes at an image quality cost. Even with this limitation, current 3D US systems are far from being portable and low-power.

B. Proposed Ultrasound Imager

In this work, we present our demonstrator of the first telesonography-capable US imager. In response to the telesonography requirements outlined above, our digital back-end prototype (i) supports 3D US, (ii) reconstructs good-quality images thanks to the scalable support for a high transducer channel count, e.g., 1024 which is state-of-the-art, (iii) is contained within a relatively inexpensive single off-the-shelf Kintex UltraScale KU040 FPGA [22], (iv) has an estimated power budget of 6W, (v) provides a complete digital signal processing pipeline from sampled echo signals to on-screen image over HDMI, (vi) offers two basic scan-converted orthogonal cross-sections to provide real-time feedback to the operator, (vii) supports various imaging modes including zone imaging and compound imaging.

It is to be noted that in principle, a device for telesonography could just be an acquisition apparatus that uploads raw data to the hospital without much processing at all, and potentially not even local image reconstruction. This however is impractical because the inexperienced operator would have no feedback whatsoever. Additionally, the amount of data to be uploaded would be much larger than the size of a fully-processed image. Similarly, the enhancements of point (vi) could be performed on the hospital side, but with a major data volume penalty, which could be a roadblock in poorly connected environments. Therefore, we focus on a prototype capable of full, high-quality autonomous image reconstruction.

This work builds on top of previous publications by our group [10]–[12], [23]. The main contributions of this paper are: (i) a thorough discussion of the zone and compound imaging features, (ii) a detailed image quality assessment including contrast ratio and peak-signal-to-noise ratio, (iii) several platform optimizations to reduce resource utilization, and (iv) additional platform configurability, allowing e.g., runtime changes of image brightness, contrast, resolution, etc.

The proposed platform supports real-time inputs from an ultrasound probe via an optical connection, as well as offline simulation inputs over an Ethernet cable. The design is highly scalable for various probe dimensions (i.e., array or matrix

probes for 2D or 3D imaging, respectively), and various numbers of transducer elements.

The paper is organized as follows: a brief literature review of recent commercial and research work is presented in Section II. A full description of the proposed system including the optimization of each processing step and the corresponding FPGA architecture is shown in Section III. This is followed by the experimental results, in Section IV, where we evaluate our imager in terms of the area and power utilization, performance, and reconstruction quality. Different quality assessment metrics have been used to evaluate the reconstruction quality, and the impact of each supported imaging mode. Finally, we summarize our conclusions and future development outlook.

II. PREVIOUS WORK

3D US imaging systems, either commercial or research-oriented, should in principle support a high input channel count, because a fully-populated, high-density matrix array can easily comprise thousands of elements [2]. However, no existing system is able to process the resulting stream of data, due to a variety of limitations including cost. This leads to quality trade-offs.

Typically, the channel count is significantly reduced, e.g., to a few hundreds, already within the piezoelectric probe. This can be achieved by different techniques. Analog pre-BF [14]–[17] is the most common; given a fully-populated transducer, it adds up the signals received by a group of transducer elements, according to a fixed delay profile, and connects the output to a single channel towards the imager. This reduces the wiring and computation complexity, but degrades the ability to focus and therefore impacts resolution.

An alternative approach is to use arrays that intrinsically feature fewer elements, or that deliberately choose to limit the number of accessible elements. For example, sparse 2D-arrays [20], [21], [24] and multiplexing [18], [19] choose the count and pattern of active receive elements per scan. Again, the main drawback is the loss of information and the reduced image quality. *Row-Column Addressed (RCA)*- or *Top-Orthogonal-to-Bottom-Electrode (TOBE)*-addressed arrays access an $N \times N$ matrix with just $2N$ signals, along the edges, instead of N^2 channels. This is achieved by accessing all the row elements together and all the column elements together, with an intrinsic averaging effect. RCA drastically reduces the processing computation cost and the wiring challenge, but again at the cost of a more limited *Field-Of-View (FOV)*, restricted focusing capability with higher sidelobes, and in general a reduced overall reconstruction quality [25]–[29].

Even though channel-reduction techniques are ubiquitous and effective, their reconstruction quality cost means that support for a higher channel count is still a desirable property. In this paper, we describe a system that can process data from up to 1024 channels, which is state-of-the-art. Our architecture could further be used in combination with one of the techniques described above to be compatible with transducers with even higher element counts.

Despite using channel-reduction techniques, current 3D US systems still end up bulky and power-consuming. On the commercial side, [1]–[3] are very advanced and provide full 3D support, but they are aimed at hospitals: they come on a base with wheels, they consume hundreds of watts and need to be permanently plugged to a power outlet, and are very expensive. Many commercial US machines that *are* portable, like the GE Voluson i [30], Samsung UGEO HM70A [31], [32], and Chison Q9 [33], are essentially 2D imagers; they nominally offer a 3D feature by supporting mechanically-swept arrays with low channel count (e.g., 128 elements). This type of probes introduces motion artifacts, and is only suitable for applications like obstetrics, where the subject (the fetus) is mostly still. Applications like cardiac imaging require a high frame rate and high resolution at the same time, and therefore full 3D support. For example, the Philips CX50 [34] scanner has a mode for 3D *Trans-Esophageal Echocardiography (TEE)*, which is considered as an invasive examination that needs not only the presence of a sonographer but also a physician, leveraging a matrix probe of 2500 elements, but analog pre-BF is exploited to compress this data onto only 100 receive channels [35]. Cephasonics provides the first US system supporting up to 4096 channels, called cQuest Griffin [36], [37], by stacking 64×64 -channel cQEngine modules. A 1024-channel cQuest Griffin consists of 16 cQEngine modules, consumes 640 W power, and has a size of $30 \text{ cm} \times 68 \text{ cm} \times 48 \text{ cm}$.

A few 3D US research systems have been proposed. They all reduce the receive channel count, yet still ending up bulky. For example, the recent second-release ULA-OP system [38] is able to perform 256-channel BF on 8 high-end FPGAs and 16 DSPs. The advanced research platform SARUS [39] supports 1024 receive channels, the highest count supported by any single 3D system, but requires 320 FPGAs. The Sonic Millip3De system [40], [41], that performs ultra-fast imaging, uses 128×96 probe elements - but only 1024 channels are considered per shot - with a powerful die-stacked package. However, its main bottleneck is the required external DRAM memory to store the BF delay coefficients, with the need of several GB/s memory bandwidth. There are also few single-FPGA US research systems, but they support up to only 64 receive channels [42], [43]. Moreover, there are many US systems relying on software-based BF, running on GPU, CPU, or DSP, but the resulting power budget is not optimal for battery operation, specially in 3D imaging. The DiPhas system [44], which supports up to 256-channel GPU-based processing and can be scaled by connecting multiple cabinets, also needs to be plugged in and is expensive.

A smart and efficient BF technique is needed to achieve a compact design, yet high-quality reconstruction. Each processing step in the imager should be optimized to yield overall platform compactness (ideally, single-chip) while being compatible with battery operation. In this work, we propose a novel, fully-scalable, complete, and single-FPGA US back-end system. We built our system around our previously proposed BF technique [10]–[12], [23]. In this design, an efficient and optimized method for each imaging step has been used to achieve overall platform efficiency, portability, and low power consumption.

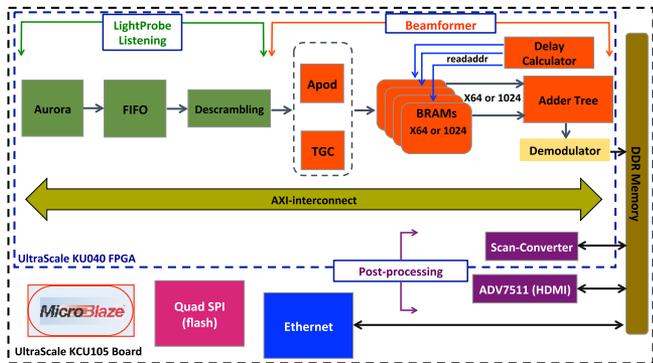


Fig. 1. Block diagram of the proposed US Imager.

TABLE I
SYSTEM SPECIFICATIONS

| Parameter | Value |
|--|--|
| Speed of sound in tissue | 1540 m/s |
| Transducer center frequency | 4 MHz |
| Transducer bandwidth | 4 MHz |
| Transducer array size (2D) | 64 elements |
| Transducer matrix size (3D) | 32×32 elements |
| Wavelength | 0.385 mm |
| Sampling frequency | 20 MHz |
| Focal points (2D) | $64 \times 500 = 32k$ |
| Imaging cross-section ($\theta \times r$, 2D) | $73^\circ \times 260\lambda$ |
| Focal points (3D) | $64 \times 64 \times 500 = 2M$ |
| Imaging volume ($\theta \times \phi \times r$, 3D) | $73^\circ \times 73^\circ \times 260\lambda$ |

III. PROPOSED US IMAGER SYSTEM

In order to achieve a single-chip US processing platform that is able to support a high number of receive channels, in particular for 3D imaging, each block in the imaging pipeline must be optimized, while preserving the image quality necessary for medical applications. We managed to fit the 3D US processing logic for up to 1024 channels into a single latest-generation but mid-range FPGA, a Kintex UltraScale KU040 [22]. Fig. 1 shows the architectural diagram of the proposed design. The main blocks of the implementation are a hardware beamformer, a hardware scan converter, and control software running on a MicroBlaze.

A. Objective and System Specifications

Our US imager architecture is highly scalable and adaptable, either statically or dynamically, for different system settings, e.g., the number of receive channels, the center frequency of the probe and its sampling frequency, the reconstruction volume and resolution, etc. However, in the following we will report numbers based on the settings specified in Table I. We present two versions of our design: a 2D US imager that supports 64 receive channels and reconstructs 64×500 focal points per frame, and a 3D US imager that supports 1024 receive channels and $64 \times 64 \times 500$ focal points. The architecture supports the full digital Rx processing pipeline of the imager, i.e., from US echoes to images, which is by far its most demanding portion.

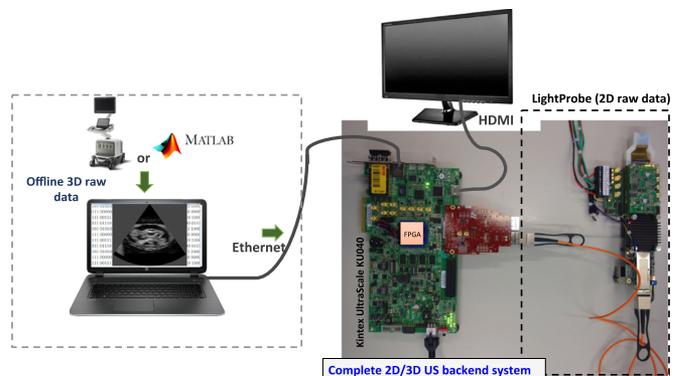


Fig. 2. Setup of the design for 2D and 3D US imaging. The components in the dashed box to the left are for development, debugging and verification only. For the probe on the right, which is not covered in this paper, refer to [45].

B. Main Imager Architecture

1) *Optical Probe Connection*: In order to provide realtime data to the processing platform, an efficient communication channel with a probe should be chosen. Due to the high bandwidth requirements, especially for 3D imaging, the two preferred options to connect a probe to our KCU105 board [46] are either over the optical SFP ports or the *Peripheral Component Interconnect Express (PCIe)* interface. The latter implies higher complexity, in terms of design and cabling, as well as higher power consumption, so the optical interface is favoured.

As a demonstration of this interface, in a collaboration with the *Swiss Federal Institute of Technology in Zurich (ETHZ)*, we have connected to our design a probe for 2D imaging. This probe comprises 64 elements and features a 4 MHz center frequency, and is designed with optical connection support [45]. The ADCs in the probe sample the data with 12-bit resolution at 20 MHz. The sustained output bandwidth is thus of around 15.3 Gbps. To support it, the *Quad enhanced Small Form-factor Pluggable (QSFP+)* interface is used by the probe; each of its four SFP+ lanes has a nominal bandwidth of 6.25 Gbps and a net one of 5 Gbps, for a total of 20 Gbps. Each lane is driven by a Xilinx Aurora 8b10b LogiCORE IP [47] for lightweight, scalable, and high-speed serial communication.

Within our Kintex FPGA, we have also implemented a 4-lane Aurora input channel. Since the KCU105 board [46] natively only features two \times SFP+ ports, we have connected an intermediate QSFP+ board to the *FPGA Mezzanine Card (FMC)* interface [48] on the KCU105 (Fig. 2).

2) *Ethernet Support For Simulated 3D Data Input*: To feed our imager with 3D data for development and debug, offline simulated data have been used. These data were obtained from the Field-II simulator [49] for Matlab fed with a description of the chosen matrix probes and 3D phantoms. For this 3D development mode, we have utilized an Ethernet port as the data transmission mean between a laptop with Matlab and the FPGA (Fig. 2). Obviously the bandwidth of the Ethernet link proves to be a performance bottleneck, but an advantage of this setup is the extra scope for flexibility and debug (Section III-G).

3) *Static Apodization and Time-Gain Compensation*: Both apodization and time-gain compensation are weighting

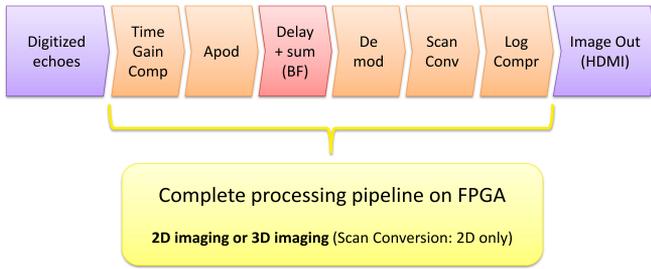


Fig. 3. Logical processing flowchart of the US image reconstruction pipeline.

functions applied to the echoes received by the US probe (Figs. 1, 3). They can be expressed through the weighting parameter $w(t, D)$ in (1) of the BF process, as follows:

$$w(t, D) = w_{apod}(t, D)w_{tgc}(t), \forall D \in 1, \dots, N. \quad (2)$$

The apodization [50] w_{apod} is a weighting window - typically a Hanning function or similar - applied to the echoes received by each piezoelectric element D of the transducer. This compensates the antenna-like behavior of the element array, that yields sidelobes in the receive directivity. Typically, the apodization window has a width that is a function of the imaging depth (as indicated by time t), called “dynamic” or “expanding-aperture” apodization [50]. However, at a shallow depth - which is 1.6 cm with our settings - the optimal window width already reaches the whole transducer’s width. Thereafter, it becomes a static function of only D , $w_{apod}(D)$. In order to reduce the resources utilized by our imager, and since the difference between static and dynamic apodization is in the very shallow and less critical region for diagnosis, we chose to implement a fully static apodization in our imager. The static apodization coefficients, one per transducer element (therefore 64 in 2D and 1024 in 3D), can be pre-calculated and stored in a small table with 16-bit representation. The total table size is 128 B in 2D and 2 kB in 3D. This fits in a single Xilinx BRAM, the size of which is 1024 rows of 18-bits. In order to apply the apodization coefficients to the RF data on the fly, a number of DSP multipliers matching the input sample rate must be deployed. When the Ethernet input is used, this means 2 samples/clock cycle, while with Aurora it is possible to receive up to 6 samples/clock cycle. Therefore, a small number of DSP multipliers is sufficient.

The US wave propagating inside the body is exposed to an attenuation directly proportional to the traveling depth. This means that later echoes (i.e., reflected by deeper scatterers) are more attenuated than earlier echoes (i.e., from shallower scatterers). A time-based compensation needs to be applied, called *Time-Gain Compensation (TGC)*. TGC can be performed in the analog domain, for example in the transducer head, but also in the digital domain, as we implemented ($w_{tgc}(t)$). In the proposed design, the TGC coefficients are also pre-calculated and stored in a small table; a small number of DSP multipliers has been used to apply these coefficients to the apodized RF data.

Typical US echoes are sampled with 12- to 16-bit ADCs. The weighted data can thus be stored onto 18-bit BRAMs in the FPGA, one sample per line (Fig. 1). To minimize memory

pressure, in streaming fashion, only a sliding window of the echo samples is on-FPGA at any time; earlier samples are discarded as soon as possible to make space for later ones. To enable this, we reconstruct images outwards by *nappes*, i.e., like onion layers [51]; traditional approaches work on *scanlines*, i.e., like pins on a pincushion, which unfortunately requires keeping the whole dataset in memory. Conveniently, geometric calculations reveal that the minimum required thickness of the sliding window is only a few hundred samples. In 3D, where BRAM count is a bottleneck, we can therefore pack samples from two receive channels into a single BRAM (i.e., 512 samples per channel simultaneously on-chip), needing a total of 512 BRAMs for 1024 channels. In 2D, we simplify the design and improve performance by dedicating one full BRAM (1024 samples) per receive channel, for a total of 64 BRAMs.

4) *Steering-Based Delay Calculation Algorithm*: The main bottleneck and the most challenging process in the 3D US processing pipeline is delay calculation (transmit delay, $t_p(|OS|)$ plus receive delay, $t_p(|SD|)$). These delays are used as indices into the raw echo data sequence, to select which samples to sum (1) to determine the reflectivity of a given scatterer (i.e., of a focal point at location S). Tx calculation is less demanding than Rx calculation by a factor of the number of transducer elements D , since O is fixed for each insonification. In 3D imaging and according to the configuration of Table I, 2 million Tx delays need to be calculated versus 2 billion Rx delays *per volume* V . Tx delays can be calculated on-the-fly on an FPGA, but for Rx calculation, at a target reconstruction rate of e.g., 50 vps, a computation bottleneck appears.

The Tx delays are calculated using directly a Xilinx CORDIC core, as seen in Fig. 4(a). The location of each scatterer $S = (x_S, y_S, z_S)$ (as a function of the current azimuth and elevation angles θ, ϕ) and of the current emission origin $O = (x_O, y_O, z_O)$ are resolved and used as inputs. Depending on the imaging mode (Sections III-C, III-D, and III-E), the initialization of the azimuth θ and elevation ϕ pointers, as well as the emission origin, changes.

To solve the challenge of Rx delay calculation, we employ an efficient approximated algorithm [10]–[12]. It is based on the first order Taylor expansion of the square root. We simplify the calculation of the enormous number of square roots per second to the exact calculation of very few square roots along the central line-of-sight ($t_p(|RD|)$), which is also done with a Xilinx CORDIC core, and then add two correction coefficients ($c_1 = \frac{x_D \sin \theta}{c}$ and $c_2 = \frac{y_D \sin \phi \cos \theta}{c}$) to calculate the delays everywhere else (3). This can be seen as “steering” the delays of the central line-of-sight, where points R are located, to reconstruct the whole frame. Fig. 4(b) shows the implementation of the Rx delay calculator. In (3), θ and ϕ are the azimuth and elevation angles, respectively, while x_D and y_D are the positions of each element D on the transducer face. r is the radial depth of the focal point being reconstructed, and c is the speed of sound in the medium, which is typically assumed constant. The architecture works identically for 2D imaging, setting ϕ and y_D to zero. Since $\theta, \phi, x_D, y_D, r$, and c have limited and deterministic values, the two correction coefficients can not only

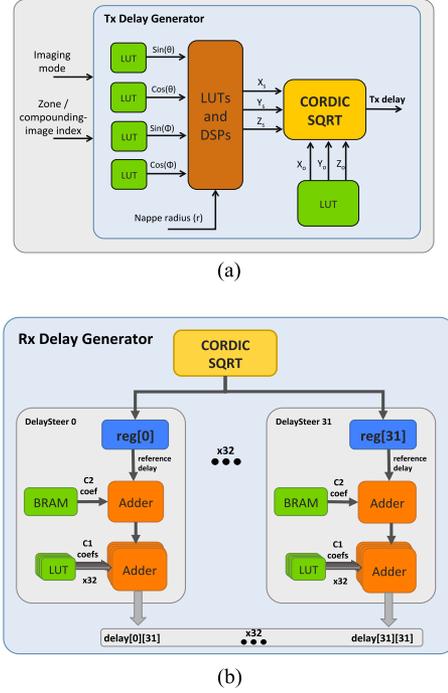


Fig. 4. The steering-based delay calculation architecture. (a) Tx delay calculator, which is based on Xilinx CORDIC IP. (b) Rx delays are computed by adding the two steering coefficients c_1 and c_2 to exactly-calculated reference delays. The sum of Tx and Rx delays is used as the access index of the input samples in the BRAMs of Fig. 1.

be calculated on-the-fly but also pre-calculated and stored in a small memory size, which we choose to do. The reconstruction rate of our beamformer is one focal point per clock cycle.

$$t_p(|S\vec{D}|) = r\sqrt{1 + \frac{x_D^2 + y_D^2}{r^2} - \frac{2x_D \sin \theta + 2y_D \sin \phi \cos \theta}{r}} \approx t_p(|R\vec{D}|) - \frac{x_D \sin \theta}{c} - \frac{y_D \sin \phi \cos \theta}{c}. \quad (3)$$

5) *Demodulation*: After BF, the reconstructed image is still in the high-frequency domain, and needs demodulation, or in other words envelope detection. Many techniques are available. A very simple demodulation technique has been implemented to reduce the resource utilization and circuit complexity. It is based on simply calculating the absolute value of the reconstructed focal points, then applying a P -order FIR low-pass filter. After each clock cycle, the absolute value of the beamformed focal point is calculated and then stored in a circular buffer of depth $P+1$ and width equal to a whole nappe. The FIR filter is applied along the lines-of-sight of the buffered nappes. Finally, the demodulated focal points are then stored in another buffer until the completion of the whole volume/image, to be then used by the visualization unit (Section III-F).

C. Single-Insonification Reconstruction Imaging (SIRI)

Our single-chip architecture supports three different US imaging modes: *Single-Insonification Reconstruction Imaging (SIRI)*, zone imaging (Section III-D), and compound imaging

(Section III-E). They are considered among the essential modes in any US imaging system.

SIRI is the most basic imaging mode. In SIRI, the reconstruction of the whole ROI is performed based on a single insonification. This insonification must use a broad beam profile, such as a diverging or plane wave. In this paper, we assume the use of a phased array transducer, therefore the natural choice is a diverging beam. To achieve this, the emission origin O is *virtually* placed behind the transducer face [12]. SIRI offers the highest possible reconstruction rate—both in terms of echo acquisition and processing—and has very low memory requirements. This comes at the cost of a degraded image quality due, in particular, to the lack of focusing on transmit; i.e., the emitted acoustic energy is spread too broadly. This reconstruction technique is at the basis of ultrafast imaging [52].

D. Zone Imaging

Zone sonography virtually divides the ROI into a number of non-overlapping zones. The zones are imaged in sequence, based on dedicated insonifications, each of which consisting of a more focused beam than the one required for SIRI. Zone imaging strikes a compromise between conventional high-quality US reconstruction, i.e., line-by-line with an insonification per line (which is completely impractical for 3D imaging due to the massive number of such lines), and SIRI (which has poorer resolution). This technique is used industrially by companies such as Zonare Medical Systems Inc. [53], [54].

The virtual division of the ROI can be done in any of the imaging dimensions. In this work, we slice the ROI in the azimuth and elevation directions. Each zone’s emission profile originates from a *virtual source* O that moves on an imaginary arc behind the transducer. Our imager currently supports 2, 4, 8 (2D) or 2×2 , 4×4 , 8×8 zones (3D). An increased zone count improves lateral resolution, but comes with a proportional cost in terms of acquisition time. Notably, the processing time, processing resources, and memory requirements are almost the same. There are only two main differences in the processing pipeline: the first is the usage of different sub-Tx delay tables for each zone according to the different emission origins O , while the second is an extra “stitching” step, performed just after BF, to join the zones in memory.

E. Spatial Compound Imaging

Our platform also supports compound imaging. Compound imaging has the main purpose of reducing image speckle due to the coherent nature of the US waves used for imaging. In this technique, the whole ROI is imaged multiple times, and the various frames are merged into a single better one leveraging their diversity. Although some diversity across subsequent reconstructions is always intrinsically present due to noise and subject movement, additional diversity is often intentionally introduced to improve the outcome. Commonly, the intentionally modified parameter is either the emission frequency—or the emission profile in space [13]. Further, compounding can be incoherent or coherent [55] depending if the

compounding happens after or before envelope detection, respectively. Incoherent spatial compounding, which we implemented, has been considered [56]–[58] as a promising technique for reducing the speckle and clutter effects [59] of the US images, with the ability to improve the organ delineations and boundaries. To support spatial compounding, we provide the possibility of steering the emission origin by $(0^\circ, \pm 10^\circ, \pm 20^\circ)$ in both azimuth and elevation. The beamformer, if so programmed, will automatically rotate between a configurable number of origin permutations, yielding slightly different images each time. The compounding operation itself is lightweight, and we choose to implement it in software on the platform’s MicroBlaze, although hardware acceleration would be possible too. Different operators can be used for compounding in order to formulate one final image out of the contributing reconstructions [58], [60]. Our design currently supports the typical *averaging-excluding-the-brightest-voxel* operator, the *minimum brightness voxel*, and the *maximum brightness voxel* compounding, and can be easily extended further. Compound imaging increases the acquisition time similarly to zone imaging, but also incurs extra processing time and brings extra memory requirements. This is because of the need for multiple insonifications, multiple full BF passes, and the storage of multiple full images (although this can be off-FPGA) before the compounding step can be performed.

F. Cross-Sectional Scan Conversion

The last processing step in our US imager pipeline is visualization (Fig. 3). This includes two main processes; log-compression and *Scan Conversion (SC)*. The brightness of the US images just after demodulation has a very high dynamic range, which makes them appear either completely black or completely white to the human eye. The log-compression operation maps the image onto a logarithmic greyscale, with appropriate contrast.

Another key operation is necessary since, in most US imaging methods including 3D US, the beamformed images are reconstructed in polar coordinates. SC transforms the image into the Cartesian coordinate space, to be displayable on a screen. This transformation is performed using interpolation, which also allows image scaling.

The SC process for a whole volume (i.e., 3D reconstruction) is too computationally expensive to fit in a single FPGA and run in realtime. However, recall that the prime goal of this work is to decouple image acquisition and diagnosis. This means that the full 3D SC can be performed remotely at the hospital, when needed; the local operator needs at most a 2D cross-sectional image display for guidance and feedback, e.g., on a built-in device screen. Therefore, we implement a block that performs this much simpler operation. In the baseline case of 2D imaging, this block nevertheless does scan-convert the whole image.

The architecture of our SC block allows the operator to choose, using the on-board push-buttons, which cross-section of the volume is to be scan-converted and displayed. The default displayed cross-section by the system is the middle elevation slice (i.e., middle azimuth-depth plane) of the volume. Other settings, like brightness and contrast control, are also

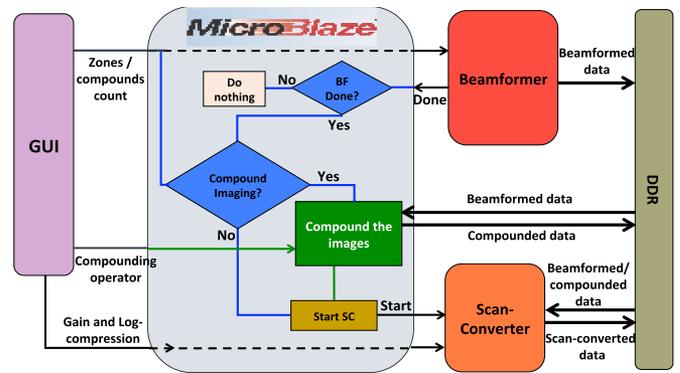


Fig. 5. The logic of controlling different imaging modes. The communication between the MicroBlaze, the FPGA, and the GUI to handle different options and settings of the processing flow.

configurable. Our SC block has been optimized over [23], [61] by equipping the IP with a master AXI interface, which bypasses the previous reliance on MicroBlaze-controlled data transfers and improves the SC rate to 14 fps. Further, the proposed design supports HDMI output to a display. We have used the Analog Devices ADV7511 [62] part on our KCU105 board for transmitting the HDMI output. The design also allows changing between different output resolutions from 640×480 to Full-HD 1920×1080 .

G. MicroBlaze, FPGA, and GUI Communication

The prototype FPGA board of this work is meant to be operable standalone. In particular, the FPGA imager is self-bootable, via a *Quad Serial Peripheral Interface (QSPI)* module that loads the boot software from the on-board flash memory.

Nevertheless, for several reasons - such as ease of configuration, debugging, ability to provide arbitrary input data, storage of the outputs, comparison to a Matlab golden reference-we provide a facility to interface the board to a GUI written in C# and running on a separate laptop. The GUI and the Microblaze control software then run in lockstep, communicating over an Ethernet cable, and the user can control all settings from the laptop (Fig. 5).

The GUI can, for example, specify “*zone count*” and “*compounding count*” parameters, which determine the imaging mode. These, in addition to others, are sent from the laptop to a packet processing stack running on the MicroBlaze; the latter then programs the BF block accordingly. Various SC settings are similarly conveyed from the GUI via the MicroBlaze to the SC block.

The board can receive the echo inputs from a probe in realtime over the optical connection (Section III-B1), with minimal involvement of the GUI and MicroBlaze. However, artificial data can also be sent from the GUI (Section III-B2), which presides over a small database of phantoms which were pre-simulated in Matlab with the Field II [49] package. In either case, the BF block then starts processing, which includes the TGC, apodization, delay-and-sum, and demodulation. The resulting image data is placed, with a direct hardware copy, onto the DDR memory off-FPGA. A status flag is then set, allowing

TABLE II
IMAGER RESOURCE UTILIZATION

| Supported Channels | Logic LUTs | Regs | BRAM | DSP | Clock | Theor. Rate |
|--------------------|------------|-------|-------|------|---------|-------------|
| 64* | 37.9% | 23.1% | 45.6% | 5.9% | 133 MHz | 4157 fps |
| 32×32* | 75.6% | 34.6% | 97.2% | 6.1% | 133 MHz | 65 vps |
| 90×90** | 79.7% | 31.5% | 69.9% | 6.5% | 133 MHz | 65 vps |

*Kintex UltraScale KU040 implementation results. **Virtex UltraScale VU190 extrapolated results.

the MicroBlaze software to detect, via polling, the completion of the process. Depending on the zone/compound imaging settings, this sequence could execute multiple times. Eventually, the MicroBlaze sends a command to the SC block, including again various parameters such as resolution, brightness, and contrast. The SC is also able to autonomously fetch the beamformed data from the external DDR, offloading the MicroBlaze from this task. The output image is quickly sent over a hardware path to the HDMI controller for display on screen, as well as stored into a different DDR location from which the MicroBlaze can - relatively slowly - send it back over Ethernet to the GUI for additional inspection.

IV. EXPERIMENTAL RESULTS

A. FPGA Utilization

In this section, we report FPGA resource utilization results achieved with Xilinx Vivado 2017.2.1. We show results for the two versions of the proposed platform: a 64-channel 2D US imager, and a 1024-channel 3D US imager. The development board we chose, which is well-suited to our requirements, is the Xilinx KCU105 with a midrange Kintex Ultrascale KU040 FPGA. Table II shows the utilization of the two proposed imagers. The two most critical resources are the BRAMs followed by the LUTs. For the BRAMs, 71% of the consumption is by the beamformer, with 42.7% (i.e., 512 BRAMs) just to store the received echoes, which is not further compressible for the given channel count. LUTs are also mostly utilized by the beamformer, followed by the platform’s AXI interconnects. Our design contains two AXI interconnects; one is configured with a “maximize performance” setting in Vivado, as necessary to support the high bandwidth in particular of the HDMI controller, while the other, which is configured with “minimize area” setting, handles the communication between the remaining blocks. The support for zone and compound imaging does not change substantially the resource utilization, except for the need of additional storage space in the external DDR to hold multiple full frames before compounding. The proposed architecture can be easily scaled up, if a larger FPGA is adopted. In the third row of Table II, we have also extrapolated the resource utilization of a high-end Virtex UltraScale XCVU190 FPGA when configuring the proposed imager to support 90×90 channels. Naturally, this chip would have much higher cost and power consumption.

The power consumption, as estimated by Vivado, of the platform is 4.6 W for the 2D configuration and 5 W for the 3D setup. This is fully aligned with the possibility of deploying the design in a battery-powered medical device. As future work, we plan to directly measure the actual power consumption of

TABLE III
IMAGER PERFORMANCE SCALING IN DIFFERENT IMAGING MODES

| Imaging Mode | Insonification Time | BF Rate | SC Rate |
|----------------------|---------------------|---------|------------------|
| SIRI | IT* | BR** | SR*** |
| Zone imaging (Z) | $Z \times IT$ | BR | SR |
| Compound imaging (C) | $C \times IT$ | BR / C | SR + Fusion time |

*Insonification time for a single emission. **BF rate for a single ROI, either volume or image, reconstruction. ***SC rate for a single frame.

the board with a dedicated Maxim tool [63]. This project will also entail a careful optimization and power management of any board components and interfaces which are not needed by the US imager.

B. Performance

In the proposed final imaging system, the acquisition rate (i.e., the cadence of the images to be stored and uploaded to the hospital) and the display rate (i.e., the images to be displayed on a local screen, after SC, for the operator to monitor) can be decoupled. The proposed imager runs at 133 MHz, and the BF rate is 1 focal point per clock cycle. This means that according to our system specifications, the *theoretical* reconstruction rate (Table II) is 4157 fps in the 2D case, and 64.9 volumes/sec (vps) in the 3D case, which is more than sufficient to upload high-temporal-resolution streams of images (movies) to a hospital. The SC rate is 14 fps at 356×300 resolution for both 2D and 3D cases. This bottleneck only applies to the real-time visualization capability, but high frame-rate reconstructions can still be performed and uploaded. The *actual* reconstruction rate also depends on the time necessary for the insonification and especially on the echo sample acquisition from the probe, which is discussed later.

The highest possible frame rate is achieved with SIRI, yielding the best frame rate at the cost of the reconstruction quality. Table III shows the performance achievable in different modes; SIRI as a reference, zone imaging of Z zones, and compound imaging of C compounds. The peak BF rate in zone imaging is the same as in SIRI since there is no overlapping of the zones - although the insonification count of Z will impact the actual reconstruction rate. Compound imaging, on the other hand, reduces both the theoretical and actual reconstruction rate linearly by a factor of C . The scan conversion cost remains constant in all modes, except that when in compounding mode, an additional averaging step of the raw frames is required.

Due to the impossibility to access a probe for 3D imaging research, a test setup requires feeding inputs and configuration settings over the Ethernet port. In this debug configuration, the output images are also transmitted back over the same connection for inspection. The Ethernet bandwidth of 1 Gbps and the related MicroBlaze IP software stack become a crucial bottleneck, resulting in a SIRI rate of about 0.02 fps for 3D. Zone and compound imaging (Table III) incur a further degradation, due to the need to send multiple insonification data sets.

A fully-functional 2D prototype with the US probe connected to the Aurora input is currently displaying output images at 0.7 fps, limited by a software bottleneck. We plan to optimize it

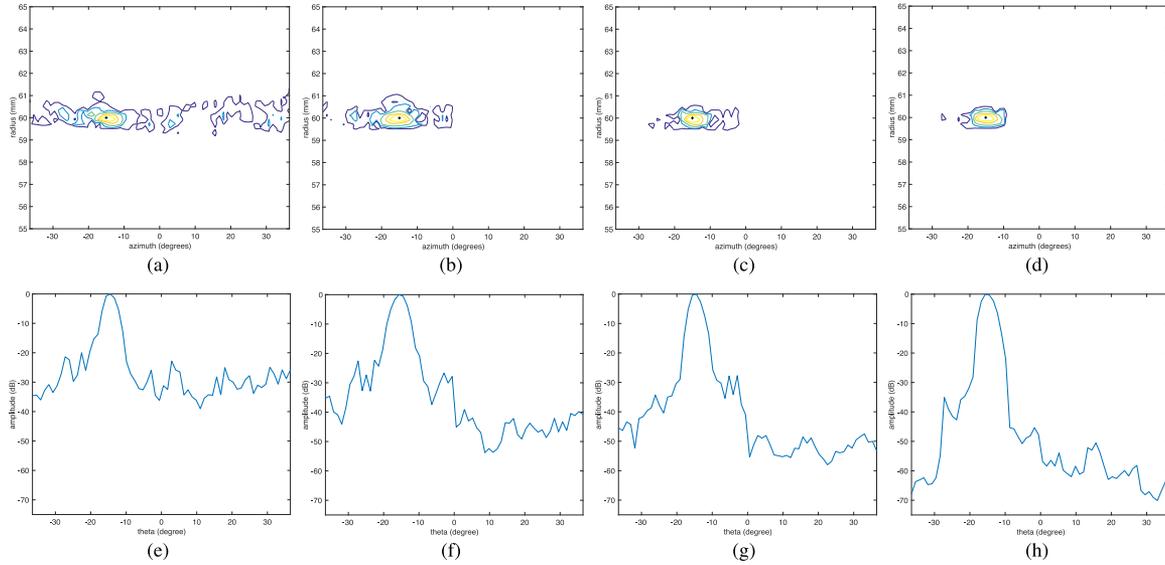


Fig. 6. Evaluation of our zone imaging implementation versus SIRI in terms of PSF contours and projections for a scatterer at $\theta = -15^\circ$, $\phi = 0^\circ$, $r = 60$ mm. This is achieved by testing the reconstruction when dividing the ROI. (a) and (e) Into SIRI. (b) and (f) Into 2×2 zones. (c) and (g) Into 4×4 zones. (d) and (h) Into 8×8 zones. The PSF contour colors represent the level in dB, in a range of -5 dB, -10 dB, -20 dB, -30 dB, -40 dB from inside to outside.

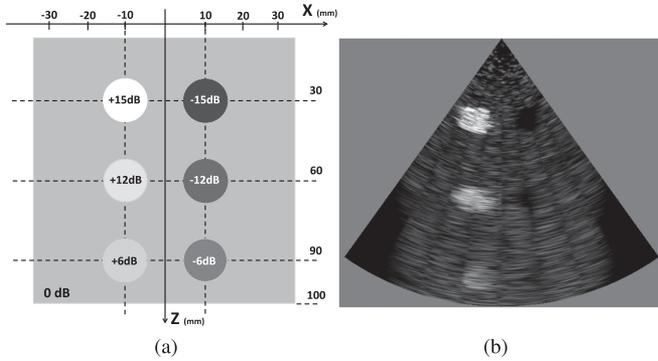


Fig. 7. Our 6-cyst phantom. (a) Sketch depicting the geometry of the phantom. (b) Phantom reconstruction using linear imaging based on a plane wave insonification.

TABLE IV
PSNR FOR COMPOUND IMAGING AND SIRI

| Image under test | PSNR (dB) |
|---|-----------|
| Single volume (SIRI)* | 34.1 |
| 3-compounding, average* | 38.8 |
| 5-compounding, average* | 41.2 |
| 9-compounding, average* | 45.2 |
| 9-compounding, average-excluding-brightest* | 41.2 |

*Reference volume: 9-frame compounding using averaging operator and based on exact delay calculation BF.

away by the use of a *Direct Memory Access (DMA)*, leading to a projected frame rate close to the 14 fps mark.

C. Image Quality Assessment

Our design achieves excellent resource efficiency at the cost of a slight inaccuracy in image reconstruction. Two of the main causes are the steering-based approximate delay calculation algorithm [12], and the use of static apodization. For the former,

we have mathematically analyzed the image quality on reference images, and have confirmed [12] - as theoretically expected - that the introduced inaccuracy occurs only very close to the probe surface and at the extreme lateral edges of the ROI. These regions are usually not critical for diagnosis, and in practice the disturbance is limited to a minor speckle pattern difference. Nonetheless, a detailed inaccuracy quantification using different metrics has been performed in [12], [64]. For what concerns static apodization, we have evaluated the impact of this optimization, finding that it only affects the shallowest 1.6 cm of the ROI, which overlaps with the inaccuracy introduced by the steering-based algorithm. The same considerations apply.

In this article, we assess the image reconstruction quality when using the different supported imaging modes: SIRI, zone- and compound imaging. Different quality metrics have been used for the purpose. The *Point Spread Function (PSF)* has been used to evaluate the lateral resolution benefits of zone imaging over SIRI. *Peak Signal-to-Noise Ratio (PSNR)* has been used to evaluate the speckle noise suppression of compound imaging over SIRI. The *Contrast Ratio (CR)* has been measured for all three modes. Further, a visual assessment is proposed in Section IV-C4.

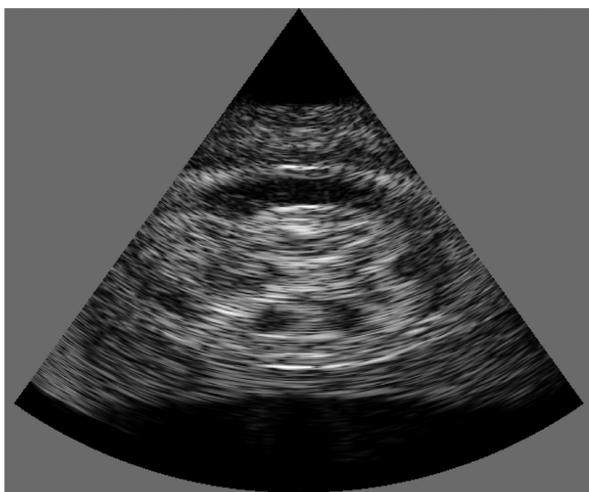
1) *Point Spread Function (PSF)*: We have evaluated our zone imaging implementation and its impact on the reconstruction quality by insonifying a single point scatterer while varying the zone count. The result is measured in terms of the PSF contours, and the projections thereof, to evaluate the lateral resolution and the height of the side lobes. A single representative scatterer at $\theta = -15^\circ$, $\phi = 0^\circ$, $r = 60$ mm has been chosen for this analysis for space reasons. We divided the ROI into 2×2 zones [Fig. 6(b) and (f)], 4×4 zones [Fig. 6(c) and (g)], and 8×8 zones [Fig. 6(d) and (h)]. The zones span a fraction of the azimuth and elevation angles, i.e., they cover sub-quasi-pyramids of the overall volume. A SIRI (i.e., 1-zone)

TABLE V
CR FOR COMPOUND IMAGING AND SIRI

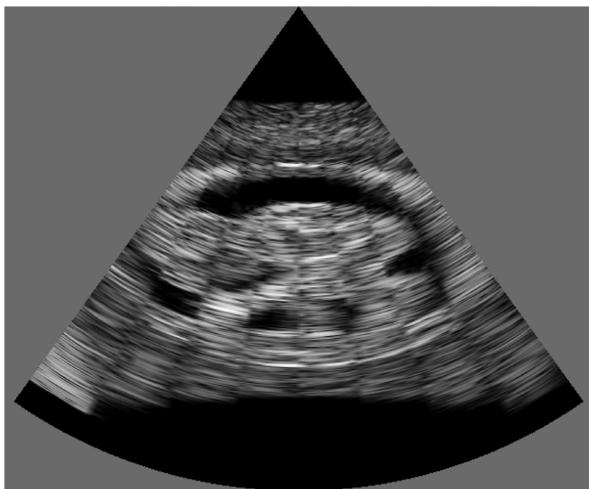
| Volume under test | +15 dB CR (dB) | +12 dB CR (dB) | +6 dB CR (dB) | -15 dB CR (dB) | -12 dB CR (dB) | -6 dB CR (dB) |
|--|-------------------|-------------------|------------------|-------------------|-------------------|------------------|
| Single volume (SIRI) | 19.3 | 16.1 | 6.8 | -3.7 | -2.6 | -2.3 |
| 3-compounding using average fusion | 18.4 | 16.7 | 7.1 | -4.9 | -2.6 | -2.2 |
| 5-compounding using average fusion | 18.2 | 16.8 | 7 | -5.1 | -2.4 | -2.3 |
| 9-compounding using average fusion | 18 | 16.5 | 7 | -5.4 | -2.7 | -2.2 |
| 9-compounding using average-excluding-brightest fusion | 18.9 | 16.6 | 7 | -5.2 | -2.1 | -1.9 |

TABLE VI
CR FOR ZONE IMAGING AND SIRI

| The volume under test | +15 dB CR (dB) | +12 dB CR (dB) | +6 dB CR (dB) | -15 dB CR (dB) | -12 dB CR (dB) | -6 dB CR (dB) |
|---------------------------|-------------------|-------------------|------------------|-------------------|-------------------|------------------|
| Single zone volume (SIRI) | 17.9 | 17 | 7.6 | -4 | -3.7 | -2.2 |
| 8 × 8 zones volume | 26.7 | 17.2 | 9.7 | -11.4 | -8.7 | -3 |



(a)



(b)

Fig. 8. 3D imaging of a fetus phantom available online on the Field II website [65]. The images show the middle azimuth-depth plane (i.e., middle elevation plane) of the reconstructed volume by our imager. The imaging volume is $73^\circ \times 73^\circ \times 10$ cm. (a) A single-zone image, i.e., SIRI, based on a single diverging beam insonification. (b) 8×8 zone image, where each zone is insonified independently with a more focused beam.

reconstruction is shown in Fig. 6(a) and (e). The expected scatterer location is denoted by a “+” symbol in the figures. For the SIRI image, the reconstruction exhibits very low lateral resolution due to distributing the insonification energy all over the ROI. Fig. 6(e) shows that the PSF projection has a wide main-lobe and a high sidelobe level. This result is expected since the insonification of a whole 3D ROI with a single US beam leads to a major dispersion of acoustic energy, which correlates negatively to resolution. The rest of the sub-figures show that there is a direct relationship between the number of zones and the improvement in the image quality, especially in terms of lateral resolution. This confirms the theoretical expectations, since the ROI is now divided in slices, each insonified with a more focused and therefore higher-energy-density beam. A notable reduction can be observed in the level of the sidelobes compared to the main-lobe, while the main-lobe width becomes narrower and the contours become more concentrated around the scatterer’s actual location.

2) *Peak Signal-to-Noise Ratio (PSNR)*: PSNR is a quality metric used to quantify how far the image under-test is from a reference image. PSNR is defined as follows:

$$\text{PSNR} = 20 \log_{10}(\text{MAX}) - 10 \log_{10}(\text{MSE}). \quad (4)$$

MAX represents the maximum voxel brightness in the image. Since we log-compress the reconstructions onto a grayscale range of (0–255), MAX is equal to 255. MSE is the Mean-Square-Error defined as:

$$\text{MSE} = \frac{1}{A \times E \times R} \sum_{a=0}^{A-1} \sum_{e=0}^{E-1} \sum_{r=0}^{R-1} [V_{ref}(a, e, r) - V_{test}(a, e, r)]^2, \quad (5)$$

where A , E , and R are the azimuth, elevation, and radial dimensions of the volume V . V_{ref} is the reference volumetric image, while V_{test} is the volumetric image under test. We have considered V_{ref} as the output volume of compounding 9 frames using the averaging operator and - most importantly - reconstructed based on *exact* delay calculation BF. The volumes

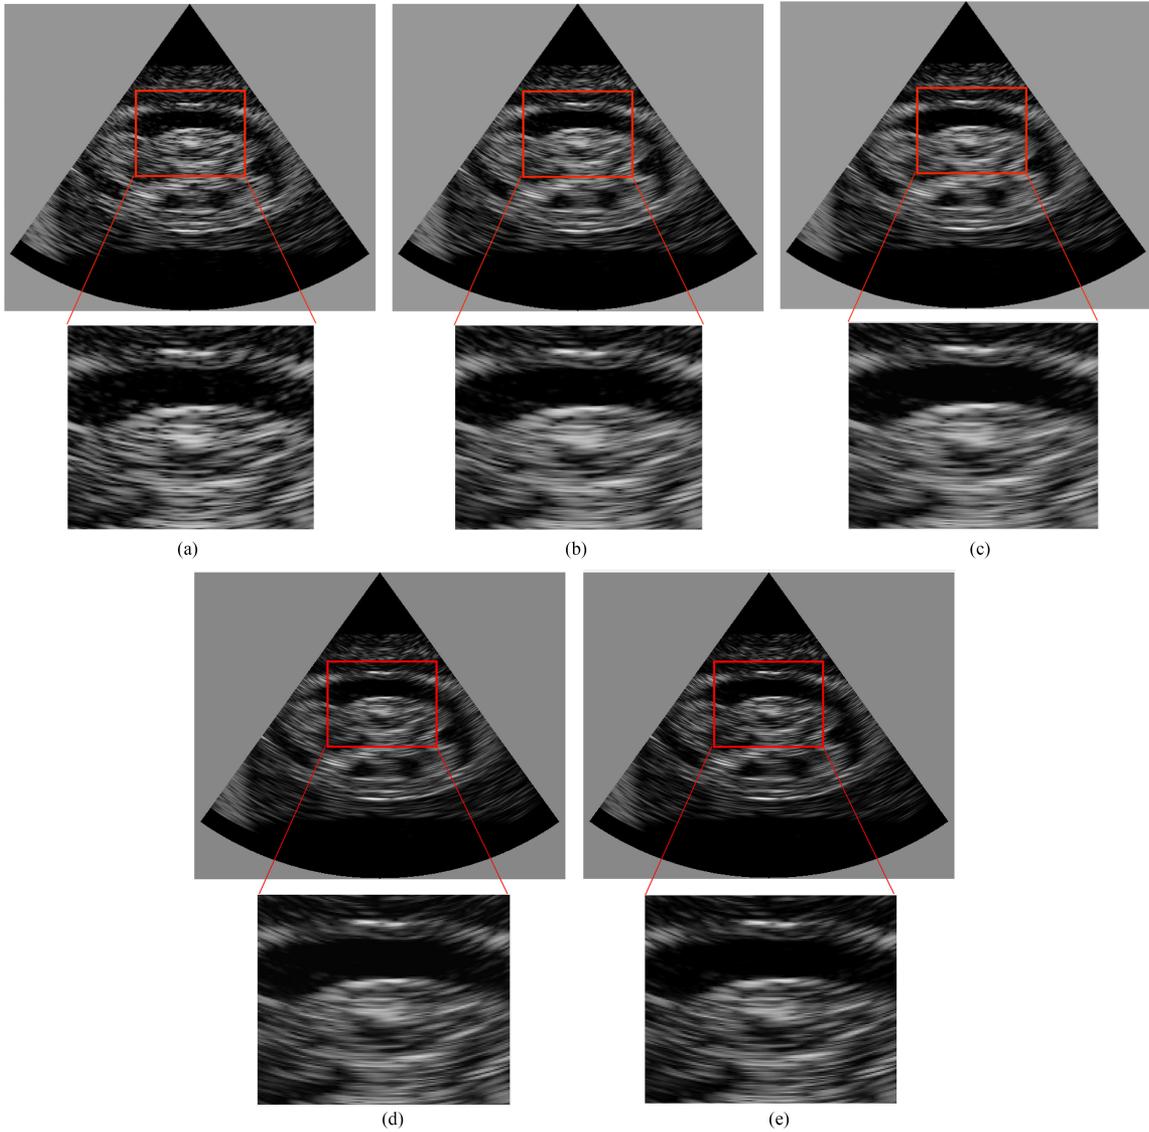


Fig. 9. 3D imaging of a fetus phantom available online on the Field II website [65]. The images show the middle azimuth-depth plane (i.e., middle elevation plane) of the volume reconstructed by our imager. The imaging volume is $73^\circ \times 73^\circ \times 10$ cm. (a) Single frame, i.e., SIRI. (b)–(d) 3-frame, 5-frame, and 9-frame compounded images, using the *average* operator. (e) 9-frame compounded image using the *average-excluding-the-brightest-voxel* operator.

under test V_{test} are the output of either SIRI or compounding different number of frames, reconstructed based on our *approximate* steered-based delay calculation approach as implemented on the FPGA. For the analysis, we have used a 6-cyst Matlab phantom developed in house and shown in Fig. 7. The phantom comprises 6 cysts of radius 5 mm, with reflectivity of +15 dB, +12 dB, +6 dB, -15 dB, -12 dB, and -6 dB compared to the background tissue.

Table IV shows the PSNR in dB for: (i) a single volume, (ii) the output of 3-, 5-, and 9-frame compounding with the *average* operator, and (iii) the output of 9-frame compounding with the *average-excluding-the-brightest-voxel* operator. The results show progressive removal of the speckle as we increase the number of compounded volumes. Moreover, the choice of the compounding operator plays an important role in the noise suppression.

3) *Contrast Ratio (CR)*: We have further appraised the image quality of our prototype outputs using the CR metric. CR is measured by selecting two regions in the image, one is for the structure of interest (i.e., a cyst), while the other is for the background. The CR is the ratio between the mean brightness of those two regions in dB, as follows:

$$CR = 20 \log_{10} \frac{\text{Mean}_{test}}{\text{Mean}_{reference}}. \quad (6)$$

For this measurement, we have used the same cyst phantom described earlier. Tables V and VI show the CR evolution when varying the compounding count and zone count, respectively. The following trends can be observed:

- The CR of each cyst does roughly follow the expected brightness trend, i.e., bright cysts are brighter than the background which is brighter than the dark cysts.

- However, the CR does not match the theoretical expectations too well in Table V and the first row of Table VI. We attribute this to the very bad lateral resolution of the imager when the whole volume is insonified with a single emission, as discussed in Section IV-C1. In practice, the bright cysts on the left side of the image “bleed into” the dark ones on the right, dramatically reducing the contrast. This effect is expected.
- *Incoherent* compound imaging is not expected to improve contrast rates significantly, and indeed has a negligible impact on CR figures (Table V).
- On the other hand, zone imaging can dramatically improve lateral resolution and therefore cancel out the bleeding effect. This is strongly confirmed by the second line of Table VI. This row shows a much better match between expected and measured CR, and a much better correlation of the relative brightness of the cysts. We expect that an even higher zone count would further improve the CR numbers, of course at a frame rate cost.

4) *Visual Assessment*: In this section, a visual assessment based on the appearance of the reconstructions is performed. We show sample 3D reconstructions from our imager based on a fetus phantom available on the Field II simulator website [65]. Fig. 8(b) shows an 8×8 multi-zone reconstruction versus a SIRI image [Fig. 8(a)]. The multi-zone reconstruction shows a clear improvement in the contrast and lateral resolution, as expected. A post-processing lateral low-pass-like filtering might be appropriate to smooth the transitions at the zone edges; however, this could be performed easily at the hospital side.

Fig. 9 shows reconstructions achieved by compounding different numbers of frames, versus a single frame. Fig. 9(a) shows single-frame reconstruction (SIRI). Fig. 9(b)–(d) show the compounding of 3, 5, and 9 frames, respectively, using a simple *average* operator. These clearly demonstrate the capability of the imager to filter out the speckle noise. Different levels of noise reduction can be achieved by adjusting the compounding degree, at the expense of the reconstruction rate. We also support different compounding operators: *average-excluding-the-brightest-voxel* [Fig. 9(e)], *maximum brightness voxel*, *minimum brightness voxel*. In the presented example, the *average-excluding-the-brightest-voxel* operator used in 9(e) shows better contrast and noise reduction than the typical *average* operator [Fig. 9(d)]. Overall, the imager provides options to tune the level of noise reduction and contrast in the final image; the best choice would depend on the structure being imaged.

V. CONCLUSIONS AND FUTURE WORK

In this work, we have developed the first single-FPGA 3D 1024-channel US back-end system. It is a fully digital system that is able to perform real-time 2D and 3D complete US reconstruction including pre-processing, BF, and post-processing. The platform supports up to 1024 receive channels, which matches the state of the art, with an unprecedented estimated power consumption of about 5 W. This has been successfully accomplished by utilizing an efficient delay calculation algorithm, in addition to optimizing each step in the processing pipeline.

The design supports both real-time inputs, over an optical connection which we could test in combination with a probe for 2D imaging, and simulated inputs, thanks to the Ethernet support for either 3D and 2D imaging. The platform natively supports image enhancement techniques like zone- and compound imaging. The proposed platform enables telemedicine by unlocking the possibility of acquisition by any untrained operator in underserved locations and situations, e.g., remote rural areas, underdeveloped regions, rescue scenarios, and battlefields.

We plan to further optimize the platform in several respects, chiefly in terms of resource occupation. We also plan to more accurately measure and optimize the power consumption and the maximum achievable frame rate. We will then focus on efficient data compression techniques for the image upload and on image navigation technology at the hospital side.

REFERENCES

- [1] GE Healthcare, “Voluson E10,” 2017. [Online]. Available: www.gehealthcare.com
- [2] Philips Healthcare, “EPIQ 7 ultrasound system,” 2015. [Online]. Available: <http://www.medical.philips.com/main/products/ultrasound/systems/epiq7/>
- [3] Samsung Electronics Co., Ltd., “Samsung WS80A with elite,” 2017. [Online]. Available: <http://www.samsung.com/global/business/healthcare/healthcare/ultrasound/US-WS8AF4K/WR>
- [4] Butterfly Network Inc., “Butterfly iq,” 2018. [Online]. Available: <https://www.butterflynetwork.com/index.html>
- [5] Philips Healthcare, “Lumify,” 2016. [Online]. Available: <https://www.lumify.philips.com/web/>
- [6] Clarius Mobile Health, “Clarius: Handheld wireless ultrasound scanner,” 2018. [Online]. Available: <https://www.clarius.me/products/>
- [7] GE Healthcare, “Vscan portfolio,” 2017. [Online]. Available: http://www3.gehealthcare.com/en/products/categories/ultrasound/vscan_portfolio
- [8] H. Hewener and S. Tretbar, “Mobile ultrafast ultrasound imaging system based on smartphone and tablet devices,” in *Proc. IEEE Int. Ultrason. Symp.*, Nov. 2015, pp. 1–4.
- [9] J. Kang *et al.*, “A system-on-chip solution for point-of-care ultrasound imaging systems: Architecture and ASIC implementation,” *IEEE Trans. Biomed. Circuits Syst.*, vol. 10, no. 2, pp. 412–423, Apr. 2016.
- [10] A. Ibrahim *et al.*, “Tackling the bottleneck of delay tables in 3D ultrasound imaging,” in *Proc. Design Autom. Test Eur. Conf.*, Mar. 2015, pp. 1683–1688.
- [11] F. Angiolini *et al.*, “1024-Channel 3D ultrasound digital beamformer in a single 5W FPGA,” in *Proc. Design Autom. Test Eur. Conf.*, Mar. 2017.
- [12] A. Ibrahim *et al.*, “Efficient sample delay calculation for 2D and 3D ultrasound imaging,” *IEEE Trans. Biomed. Circuits Syst.*, vol. 11, no. 4, pp. 815–831, Aug. 2017.
- [13] U. Bencharit, J. L. Kaufman, N. M. Bilgutay, and J. Saniie, “Frequency and spatial compounding techniques for improved ultrasonic imaging,” in *Proc. IEEE Ultrason. Symp.*, Nov. 1986, pp. 1021–1026.
- [14] J. Larson, “2-D phased array ultrasound imaging system with distributed phasing,” U.S. Patent 5,229,933, Jul. 20, 1993. [Online]. Available: <https://www.google.ch/patents/US5229933>
- [15] B. Savord and R. Solomon, “Fully sampled matrix transducer for real time 3D ultrasonic imaging,” in *Proc. IEEE Symp. Ultrason.*, 2003, vol. 1, pp. 945–953.
- [16] J.-Y. Um *et al.*, “A single-chip 32-channel analog beamformer with 4-ns delay resolution and 768-ns maximum delay range for ultrasound medical imaging with a linear array transducer,” *IEEE Trans. Biomed. Circuits Syst.*, vol. 9, no. 1, pp. 138–151, Feb. 2015.
- [17] H. g. Kang *et al.*, “Column-based micro-beamformer for improved 2D beamforming using a matrix array transducer,” in *Proc. Biomed. Circuits Syst. Conf.*, 2015.
- [18] T. M. Carpenter, M. W. Rashid, M. Ghovanloo, D. Cowell, S. Freear, and F. L. Degertekin, “Time-division multiplexing for cable reduction in ultrasound imaging catheters,” in *Proc. Biomed. Circuits Syst. Conf.*, 2015.
- [19] J. T. Yen and S. W. Smith, “Real-time rectilinear 3-D ultrasound using receive mode multiplexing,” *IEEE Trans. Ultrason., Ferroelectr., Freq. Control.*, vol. 51, no. 2, pp. 216–226, Feb. 2004.

- [20] A. Austeng and S. Holm, "Sparse 2-D arrays for 3-D phased array imaging—Design methods," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control.*, vol. 49, no. 8, pp. 1073–1086, Aug. 2002.
- [21] J. T. Yen, J. P. Steinberg, and S. W. Smith, "Sparse 2-D array design for real time rectilinear volumetric imaging," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control.*, vol. 47, no. 1, pp. 93–110, Jan. 2000.
- [22] Xilinx Inc., "Ultrascale FPGA: Product tables and product selection guide," 2016. [Online]. Available: <http://www.xilinx.com/support/documentation/selection-guides/ultrascale-fpga-product-selection-guide.pdf#KU>
- [23] A. Ibrahim *et al.*, "Inexpensive 1024-Channel 3D teleosonography system on FPGA," in *Proc. IEEE Biomed. Circuits Syst. Conf.*, 2017.
- [24] J. Zhou *et al.*, "High volume rate 3D ultrasound imaging based on synthetic aperture sequential beamforming," in *Proc. IEEE Int. Ultrason. Symp.*, Sep. 2017, pp. 1–4.
- [25] C. H. Seo and J. T. Yen, "A 256×256 2-D array transducer with row-column addressing for 3-D rectilinear imaging," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control.*, vol. 56, no. 4, pp. 837–847, Apr. 2009.
- [26] C. E. M. Demore, A. W. Joyce, K. Wall, and G. R. Lockwood, "Real-time volume imaging using a crossed electrode array," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control.*, vol. 56, no. 6, pp. 1252–1261, Jun. 2009.
- [27] C. Ceroici, T. Harrison, and R. J. Zemp, "Fast orthogonal row-column electronic scanning with top-orthogonal-to-bottom electrode arrays," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control.*, vol. 64, no. 6, pp. 1009–1014, Jun. 2017.
- [28] H. Bouzari, M. Engholm, M. B. Stuart, E. V. Thomsen, and J. A. Jensen, "Improved focusing method for 3-D imaging using row-column-addressed 2-D arrays," in *Proc. IEEE Int. Ultrason. Symp.*, Sep. 2017, pp. 1–4.
- [29] H. Bouzari *et al.*, "3-D imaging using row-column-addressed 2-D arrays with a diverging lens: Phantom study," in *Proc. IEEE Int. Ultrason. Symp.*, Sep. 2017, pp. 1–4.
- [30] GE Healthcare, "Voluson i," 2017. [Online]. Available: http://www3.gehealthcare.co.uk/en-gb/products/categories/ultrasound/voluson/voluson_i
- [31] Samsung Healthcare, "HM70A with plus," 2017. [Online]. Available: <http://www.samsung.com/global/business/healthcare/healthcare/ultrasound/USS-HM7AF1K/WR>
- [32] kpi Healthcare, "Samsung UGEO HM70A ultrasound machine," 2017. [Online]. Available: <http://www.kpiultrasound.com/medical-equipment/ultrasound-scanners/samsung-ultrasound/samsung-ugeo-hm70a-portable/>
- [33] Chison Medical Imaging Co., Ltd., "Chison q9," 2014. [Online]. Available: http://www.chison.com/products_detail.asp?id=168
- [34] Philips Healthcare, "Cx50 xMATRIX," 2017. [Online]. Available: <http://www.usa.philips.com/healthcare/product/HC795076CA/cx50-xmatrix-c-ardiology-ultrasound-machine>
- [35] G. Wayne Moore, "Ultrasound beamformation and image reconstruction," in *Proc. 51st Amer. Assoc. Physicists Med. Ann. Meeting*, 2009. [Online]. Available: <https://vimeo.com/88175266>
- [36] Cephasonics, "cQuest Griffin—Product Brief," Catalog, 2015. [Online]. Available: http://pdf.medicaexpo.com/pdf/cephasonics/cquest-griffin/92255-171445-_2.html
- [37] R. J. Tobias and M. E. Schafer, "System development improvements using stackable high-channel-count ultrasound hardware," in *Proc. IEEE Int. Ultrason. Symp.*, Sep. 2014.
- [38] E. Boni *et al.*, "ULA-OP 256: A 256-channel open scanner for development and real-time implementation of new ultrasound methods," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control.*, vol. 63, no. 10, pp. 1488–1495, Oct. 2016.
- [39] J. Jensen *et al.*, "Sarus: A synthetic aperture real-time ultrasound system," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control.*, vol. 60, no. 9, pp. 1838–1852, Sep. 2013.
- [40] R. Sampson, M. Yang, S. Wei, C. Chakrabarti, and T. Wensch, "Sonic millipede: A massively parallel 3D-stacked accelerator for 3D ultrasound," in *Proc. IEEE 19th Int. Symp., High Performance Comput. Architecture*, Feb. 2013, pp. 318–329.
- [41] R. Sampson *et al.*, "FPGA implementation of low-power 3D ultrasound beamformer," in *Proc. IEEE Int. Ultrason. Symp.*, Nov. 2015.
- [42] M. Tsakalakis and N. G. Bourbakis, "Designing of a low-cost, volumetric multi-transducer phased array ultrasound system," in *Proc. IEEE 15th Int. Conf. Bioinform. Bioeng.*, Nov. 2015, pp. 1–7.
- [43] Y.-J. Kim *et al.*, "A single-chip 64-channel ultrasound RX-beamformer including analog front-end and an LUT for non-uniform ADC-sample-clock generation," *IEEE Trans. Biomed. Circuits Syst.*, vol. 11, no. 1, pp. 87–97, Feb. 2017.
- [44] C. Risser, H. J. Welsch, H. Fonfara, H. Hewener, and S. Tretbar, "High channel count ultrasound beamformer system with external multiplexer support for ultrafast 3D/4D ultrasound," in *Proc. Int. Ultrason. Symp.*, Sep. 2016, pp. 1–4.
- [45] P. A. Hager, C. Risser, P.-K. Weber, and L. Benini, "LightProbe: A 64-channel programmable ultrasound transducer head with an integrated front-end and a 26.4 Gb/s optical link," in *Proc. IEEE Int. Symp. Circuits Syst.*, 2017.
- [46] Xilinx Inc., "Xilinx kintex ultraScale FPGA KCU105 evaluation kit," 2017. [Online]. Available: <https://www.xilinx.com/products/boards-and-kits/kcu105.html>
- [47] Xilinx Inc., "Aurora 8B/10B," 2017. [Online]. Available: <https://www.xilinx.com/products/intellectual-property/aurora8b10b.html>
- [48] HiTech Global LLC, "2-port QSFP+ (40G or 56G) FMC module (vita57.1)," 2017. [Online]. Available: http://www.hitechglobal.com/FMCModules/FMC_2QSFP+.htm
- [49] J. Jensen, "Field: A program for simulating ultrasound systems," in *Proc. 10th Nordic-Baltic Conf. Biomed. Imaging Published Medical Biol. Eng. Comput.*, Suppl. 1, vol. 34, 1996, pp. 351–353.
- [50] C. Daft and W. Engeler, "Windowing of wide-band ultrasound transducers," in *Proc. IEEE Ultrason. Symp.*, 1996, vol. 2, pp. 1541–1544.
- [51] P. Vogel, A. Bartolini, and L. Benini, "Efficient parallel beamforming for 3D ultrasound imaging," in *Proceedings of the 24th Edition of the Great Lakes Symposium on VLSI (GLSVLSI '14)*. New York, NY, USA: ACM, 2014, pp. 175–180. [Online]. Available: <http://doi.acm.org/10.1145/2591513.2591599>
- [52] M. Tanter and M. Fink, "Ultrafast imaging in biomedical ultrasound," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control.*, vol. 61, no. 1, pp. 102–119, Jan. 2014.
- [53] D. Napolitano *et al.*, "Continuous transmit focusing method and apparatus for ultrasound imaging system," US Patent 8,002,705, Aug. 23 2011. [Online]. Available: <https://www.google.com/patents/US8002705>
- [54] G. McLaughlin, "Zone sonography: What it is and how it's different," ZONARE Medical Systems, Inc., Tech. Rep., 2012. [Online]. Available: www.gehealthcare.com
- [55] G. Montaldo, M. Tanter, J. Bercoff, N. Benech, and M. Fink, "Coherent plane-wave compounding for very high frame rate ultrasonography and transient elastography," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control.*, vol. 56, no. 3, pp. 489–506, Mar. 2009.
- [56] S. K. Jespersen, J. E. Wilhjelm, H. Sillesen, "Multi-angle compound imaging," *Ultrasonic Imaging*, vol. 20, no. 2, pp. 81–102, Apr. 1998.
- [57] R. R. Entekin, B. A. Porter, H. H. Sillesen, A. D. Wong, P. L. Cooperberg, and C. H. Fix, "Real-time spatial compound imaging: Application to breast, vascular, and musculoskeletal ultrasound," *Seminars Ultrasound, CT MRI*, vol. 22, no. 1, pp. 50–64, 2001. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0887217101900186>
- [58] A. Perperidis, "Postprocessing approaches for the improvement of cardiac ultrasound b-mode images: A review," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control.*, vol. 63, no. 3, pp. 470–485, Mar. 2016.
- [59] V. Chan and A. Perlas, "Basics of ultrasound imaging," in *Atlas of Ultrasound-Guided Procedures in Interventional Pain Management*, S. N. Narouze, Ed. New York, NY, USA: Springer-Verlag, 2011, ch. 2, pp. 13–19.
- [60] J. E. Wilhjelm, M. S. Jensen, S. K. Jespersen, B. Sahl, and E. Falk, "Visual and quantitative evaluation of selected image combination schemes in ultrasound spatial compound scanning," *IEEE Trans. Medical Imaging*, vol. 23, no. 2, pp. 181–190, Feb. 2004.
- [61] A. Ibrahim *et al.*, "Single-FPGA complete 3D and 2D medical ultrasound imager," in *Proc. Conf. Design Architectures Signal Image Process.*, 2017.
- [62] Analog Devices Inc., "ADV7511: 225 MHz, high performance HDMI transmitter with arc," 2017. [Online]. Available: <http://www.analog.com/en/products/audio-video/analoghdmidvi-interfaces/hdmidvi-transmitters/adv7511.html#product-overview>
- [63] Xilinx Inc., "Best practices for ultraScale power measurement using the KCU105," [Online]. Available: <https://www.xilinx.com/video/fpga/best-practices-ultrascale-power-kcu105.html>
- [64] A. Ibrahim, F. Angiolini, M. Arditi, J.-P. Thiran, and G. De Micheli, "Apodization scheme for hardware-efficient beamformer," in *Proc. 12th Conf. Ph.D. Res. Microelectronics Electron.*, 2016.
- [65] "Field II simulation program: Calculation of b-mode image of synthetic fetus," Field II, Apr. 2012. [Online]. Available: http://field-ii.dk/?examples/fetus_example/fetus_example.html



and PRIME 2016 Conference, respectively.



Shuping Zhang received the B.E. degree from the Beijing Institute of Technology, Beijing, China, in 2011, and the M.E. degree from Waseda University, Kitakyushu, Japan, in 2014. He is currently working toward the Ph.D. degree at Waseda University. His current research interests include 3-D integration implementations for video coding applications.



Federico Angiolini received the M.S. degree (*summa cum laude*) in electrical engineering from the University of Bologna, Bologna, Italy, in 2003, and the Ph.D. degree from the Department of Electronics and Computer Science, University of Bologna, in 2008. His initial research interests included memory hierarchies, multiprocessor-embedded systems, and networks-on-chip, and resulted in him co-founding and being CEO of iNoCs Structured Interconnects, based in Lausanne, Switzerland. Since 2013, he has been with the Swiss Federal Polytechnical School of

Lausanne, where he is working on medical imaging and drug delivery platforms.



Marcel Arditi (M'82) was born in Alexandria, Egypt, in 1951. He received the M.Sc. degree in physics from the University of Geneva, Geneva, Switzerland, in 1975 and the Ph.D. degree from the Department of Medical Biophysics, University of Toronto, Toronto, ON, Canada, in 1982, specializing in the use of annular arrays in medical ultrasound imaging. From 1982 to 1985, he was with SRI International, Menlo Park, CA, USA, where he conducted research for various industrial and government organizations, still in the ultrasound imaging field. In

1985, he joined Battelle-Europe, Geneva, Switzerland, and managed contract research projects for corporations in the optical and automotive industries, in fields related to image and signal processing. Between 1994 and 2014, he was a Senior Scientist with Bracco Suisse SA, Geneva; he contributed to the acoustical specifications and characterization of microbubble-based contrast agents for medical ultrasound and developed software for quantifying blood perfusion with contrast-enhanced ultrasound. Since 2015, he has been with Swiss Federal Institute of Technology, Lausanne, Switzerland, working on novel beamforming strategies for medical ultrasound.



Shinji Kimura received the B.E., M.E., and D.Eng. degrees in information science from Kyoto University, Kyoto, Japan, in 1982, 1984, and 1989, respectively. He was a Visiting Scientist with Carnegie Mellon University, Pittsburgh, PA, USA, from 1989 to 1990, and a visiting scholar with Stanford University, Stanford, CA, USA, from 2000 to 2001. He has been an Assistant Professor with Kobe University, Kobe, Japan, since 1985, and an Associate Professor with the Nara Institute of Science and Technology, Ikoma, Japan, since 1993. He has been a Professor

with Waseda University, Kitakyushu, Japan, since 2002. His current research interests include the formal and timing verification of logic circuits, the hardware/software co-design methodologies, reconfigurable hardware, and the low-power design. He is a Fellow of IEICE, and a member of the Information Processing Society of Japan and the IEEE Computer Society. He was an Executive Committee Member of ICCAD 2011 and 2012, and the General Chair of ASP-DAC 2013.



Satoshi Goto (LF'11) received the B.E. and the M.E. degrees in electronics and communication engineering and the D.Eng. degree from Waseda University, Kitakyushu, Japan, in 1968, 1970, and 1978, respectively. In 1970, he joined NEC Laboratories, where he worked for LSI design, Multimedia system and Software as GM and Vice President. From 2002 to 2015, he was a Professor with the Graduate school of Information, Production and Systems, Waseda University, and is currently an Emeritus Professor with Waseda University and the Vice President of E-JUST,

Egypt. He was the GC of ICCAD, ASPDAC, VLSI-SOC, ASICON, and ISOCC and was a board member of IEEE CAS Society. He is an IEICE Fellow. He was a recipient of a number of awards and honors, including the Distinguished Achievement Awards from IEICE, the Best paper award from ISSCC, IEEE VLSI Symposium, and IEICE, and also the Jubilee Medal from IEEE and the Life Time Achievement Award from IEEE ISPD. He is Visiting Professor with Shanghai Jiao Tang University, Sun Yat-sen University, and Tsinghua University of China and a member of Science Council of Japan.



Jean-Philippe Thiran was born in Namur, Belgium. He received the degree in electrical engineering and the Ph.D. degree from the Université catholique de Louvain, Louvain-la-Neuve, Belgium, in 1993 and 1997, respectively. He joined the Swiss Federal Institute of Technology (EPFL), Lausanne, Switzerland, in February 1998, as a Senior Lecturer. He is currently an Associate Professor of signal processing and the Director of the Signal Processing Lab 5, EPFL. He also holds a 20% Associate Professor position with the Department of Radiology, University Hos-

pital Center, and the University of Lausanne (CHUV-UNIL). His current scientific interests include image analysis and computer vision, medical imaging, and multimodal signal/image processing, with application to brain connectivity analysis, remote sensing, human-computer interactions, etc. He is author or coauthor of several book chapters, 150 journal papers, and more than 200 conference papers in image processing and analysis. He holds seven international patents. Until 2015, he was an Associate Editor for the IEEE TRANSACTIONS ON IMAGE PROCESSING. Among other duties, he was also the Technical Co-Chair of the 2015 IEEE International Conference on Image Processing (IEEE ICIP 2015). He was a member of the Machine Learning for Signal Processing (MLSP) Technical Committee (2008–2010) and of the Image, Video, and Multidimensional Signal Processing (IVMSP) Technical Committee (2009–2014) of the IEEE Signal Processing Society.



Giovanni De Micheli (F'94) is a Professor and the Director of the Institute of Electrical Engineering, Swiss Federal Institute of Technology, Lausanne, Switzerland. He is the program leader of the Nano-Tera.ch program. He was a Professor with the Department of Electrical Engineering, Stanford University, USA. His research interests include several aspects of design technologies for integrated circuits and systems, such as synthesis for emerging technologies, networks on chips, and 3-D integration. He is also interested in heterogeneous platform design including

electrical components and biosensors, as well as data processing of biomedical information. He is the author of book *Synthesis and Optimization of Digital Circuits* (McGraw-Hill, 1994), and coauthor and/or coeditor of eight other books and of more than 750 technical articles. He is an International Honorary Member of the American Academy of Arts and Sciences, a member of Academia Europaea, and a Fellow of ACM. He has an H-index of 93 according to Google Scholar. He is a member of the Scientific Advisory Board of IMEC (Leuven, B), CfaED (Dresden, D), and STMicroelectronics. He was the recipient of the 2016 IEEE/CS Harry Goode award for seminal contributions to design and design tools of Networks on Chips, the 2016 EDAA Lifetime Achievement Award, the 2012 IEEE/CAS Mac Van Valkenburg award, and the 2003 IEEE Emanuel Piore Award. He has been serving IEEE in several capacities, namely Division 1 Director (2008–2009), cofounder and President Elect of the IEEE Council on EDA (2005–2007), President of the IEEE CAS Society (2003), Editor-in-Chief of the IEEE TRANSACTIONS ON COMPUTER-AIDED DESIGN OF INTEGRATED CIRCUITS AND SYSTEMS (1997–2001). He has been the Chair of several conferences, including Memocode (2014), DATE (2010), pHealth (2006), VLSI SOC (2006), DAC (2000), and ICCD (1989).