

Inexpensive 1024-Channel 3D Telesonography System on FPGA

A. Ibrahim*, D. Doy*, C. Loureiro*, E. Pignat*, F. Angiolini*, M. Arditi*, J.-P. Thiran*[†], and G. De Micheli*

* École Polytechnique Fédérale de Lausanne (EPFL), Switzerland

[†]Department of Radiology, University Hospital Center (CHUV) and University of Lausanne (UNIL), Switzerland

Abstract—Volumetric ultrasound (US) is a very promising development of medical US imaging. An under-exploited advantage of volumetric US is the mitigation of the strict probe positioning constrains necessary to acquire 2D scans, potentially allowing the decoupling of US image acquisition and diagnosis. However, today’s 3D US systems are large and beset by high power and cost requirements, making them only available in well-equipped hospitals. In this study, we propose the first telesonography-capable medical imaging system that supports up to 1024 channels, on par with the state of the art. As a first embodiment, we have implemented our design in a single development FPGA board of $26.7\text{cm} \times 14\text{cm} \times 0.16\text{cm}$, with an estimated power consumption of 6.1 W. Moreover, we have equipped our platform with an automatic positioning module to help any operator defining the scan location, hence allowing for better remote diagnosis. Our design supports two types of data inputs: real-time via an optical connection and offline over Ethernet. The reconstructed images can be visualized on an HDMI screen. The estimated cost of the proposed prototype materials is less than 4000€.

I. INTRODUCTION

Medical ultrasound (US) imaging is a well established imaging technique; nonetheless, there are still limitations to its pervasiveness. US images are typically 2D, and the US probe must be operated by a well-trained sonographer to precisely locate the body slices of interest for diagnosis. In many cases, the presence of such a professional is very difficult, like in rescue situations, in developing regions, and in special environments (ships, battlefields, and aircrafts).

3D US, an emerging technology, acquires volumetric scans. In theory, the expanded field of view allows even an untrained operator to perform the scans, hence opening the door for the decoupling of acquisition from diagnosis, called *telesonography*. However, today’s volumetric US systems are hospital-grade, and therefore stationary, expensive, and very power-consuming (500W or more), requiring a constant connection to the power grid. Therefore, they remain inaccessible to places where electricity cut-off is an issue and the medical equipment budget is limited, like in underdeveloped regions and remote rural areas. At the moment, no 3D US imaging system exists that is inexpensive, portable, and battery-operated, including research systems.

In US imaging, a Radio-Frequency (RF) (2-20 MHz) acoustic wave is transmitted through the region of interest (ROI), which ideally should consist mostly of soft tissues. Tissue inhomogeneities scatter the incident waves, acting as secondary point sources. Part of the reflected echoes is detected by the transducer and then sent to a processing unit. 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). By BF, the location and reflectivity of the tissue inhomogeneities (i.e. scatterers) are defined. 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 and backscattered to a transducer element D . The beamformed image is then post-processed to make it interpretable by the human eye when displayed on a screen. By using a probe comprising a linear array of elements D , a 2D body slice can be imaged. If the probe is matricial, a quasi-pyramidal 3D volume can be scanned.

It is very challenging to fit a 3D US system in a portable, battery powered, and inexpensive embodiment. This is due to the tremendous computational load when processing data from a high

number of input channels, and the consequent material requirements. For example, hundreds of billions of square-root calculations per second (Section III-D1) are needed to compute wave propagation times.

In this work, we have developed the first *telesonography-capable* US imaging platform supporting 1024 receive channels. The proposed design architecture has been implemented in a single Kintex UltraScale KCU105 development board [1] - based on a KU040 FPGA [2] - that has a very portable size of $26.7\text{cm} \times 14\text{cm} \times 0.16\text{cm}$. The proposed imager architecture is scalable to support up to 1024 receive channels, which is state-of-the-art, with an estimated power consumption of 6.1 W. In addition, we have developed an Inertial Measurement Unit (IMU), which can be integrated in any probe head to detect its position and transmit it back to the imager. This helps the operator tag the scan locations, in turn helping the *remote* diagnosis process.

Our telesonographic platform can be deployed, as shown in Fig. 1, in rescue situations, remote under-served areas, developing regions, and special environments like aircrafts, ships, battlefields, etc.. In all of these cases, an operator, who can be a paramedic or a very lightly trained person, acquires 3D scans for the region of interest and uploads them to a hospital, where a sonographer can diagnose and archive the scan. The location of the acquisitions can be tagged either automatically by the proposed IMU, guiding the sonographer in the diagnosis, or else with one of the aforementioned techniques.

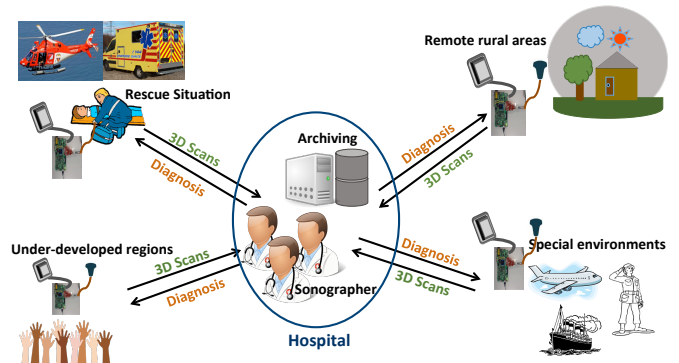


Fig. 1. Workflows enabled by the proposed telesonographic imager.

II. PREVIOUS WORK

Much research has been performed to overcome the computation bottleneck of 3D US reconstruction. However, the main solution is to reduce the number of receive channels, even if the number of transducer elements is high. This has been accomplished by different techniques, like analog pre-BF [3] and multiplexing. This comes at the cost of the image quality due to the reduced collected independent information, and the systems still end up bulky and consuming high power. For example, on the commercial side, 3D US systems [4], which mostly use analog pre-BF, are very advanced, but they can only be available in well-equipped hospitals due to their considerable space, power, and cost requirements. Semi-compact systems like the

Philips CX50 [5] can support full 3D Trans-Esophageal Echocardiography (TEE) imaging with a 2500-element matrix probe, but this is possible because analog pre-BF is used to reduce the channel count to just 100, at a major image quality cost. Cephasonics provides the first US system supporting up to 4096 channels, called cQuest Griffin [6], [7], 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 $30\text{cm} \times 68\text{cm} \times 48\text{cm}$.

On the research side, the powerful and flexible SARUS [8] system, that supports up to 1024 receive channels, runs on 320 FPGAs and needs extensive cooling. Similarly, the Sonic Millip3De [9] supports up to 1024-channel processing using a die-stacked package, but its main limitation is the required external DRAM memory for BF delays and the several GB/s memory bandwidth. The recently-released 256-channel ULA-OP system [10] is implemented with 9 high-end FPGAs and 17 DSPs with dimensions of $34\text{cm} \times 30\text{cm} \times 26\text{cm}$, for a weight of 16 kg.

Portable systems exist in the commercial space [11]–[13]. However, they are offered for 2D imaging and, at most, provide partial support for a 3D mode. In this mode, they rely on mechanically-swept transducers of few receive channels (e.g. 128). This 3D reconstruction approach requires about 1 second per frame, leading to low-quality images due to major motion artifacts. They are mainly offered for obstetrics since the fetus barely moves, and not for applications like cardiology where high resolution and frame rate are both essential.

III. PROPOSED TELESONOGRAPHY-CAPABLE IMAGER

Realizing a telesonography-capable imager with a high number of receive channels means supporting the processing of a high number of signals in a portable and battery-powered embodiment. Achieving these features is quite challenging due to the massive computation requirements of volumetric reconstruction. In this section we describe the imaging pipeline and how we tune it for a compact and efficient implementation.

A. System Specifications

Our proposed platform is highly scalable, and can either be scaled up on a larger FPGA or scaled down if only 2D imaging is required. The descriptions and experiments in this paper are based on the specifications in Table I. The proposed imager supports $32 \times 32 = 1024$ channels for 3D imaging, reconstructing 3D images of $64 \times 64 \times 500 = 2M$ voxels.

B. Signal Acquisition and Pre-processing

The design supports two input interfaces. For real-time live inputs from a piezoelectric probe, we support a high-bandwidth optical connection to our board. For testing, e.g. with simulated or pre-acquired inputs coming from a laptop, an Ethernet interface is offered.

1) *Real-Time Data Input by Optical Connection:* For this interface, we leverage a Quad Small Form-factor Pluggable (QSFP+) connection to our development board. The theoretical maximum bandwidth allowed by the SFP specification is of 16 Gbps per link, i.e. 64 Gbps for the whole QSFP+ link.

Since no open matrix probe for 3D imaging is available, as a proof of concept, we test the realtime data supply to the processing board using an array probe for 2D imaging. This probe was developed at

TABLE I
SYSTEM SPECIFICATIONS

Parameter	Value
Speed of sound in tissue	1540 m/s
Transducer center frequency	4 MHz
Transducer matrix size	32×32 elements
Sampling frequency	20 MHz
Element directivity (acceptance angle)	0.707 rad
3D Focal points	$64 \times 64 \times 500 = 2M$ voxels
Imaging volume ($\theta \times \phi \times r$)	$73^\circ \times 73^\circ \times 260\lambda$

the Integrated Systems Laboratory (IIS) of the Swiss Federal Institute of Technology in Zurich (ETHZ) [14]. It features 64 piezoelectric elements with a center frequency f_0 of 4 MHz. The echo data is fully sampled by ADCs having a resolution of 12 bits and 20 MHz. Accordingly, the required sustained bandwidth is 15.3 Gbps. Communication occurs over an Aurora 8b10b LogiCORE IP [15] on the FPGA, which adds an additional 25% bandwidth overhead for flow control and error correction. The total resulting bandwidth of about 19 Gbps is easily available on the QSFP+ connection. A detailed discussion of the interfacing is provided in [16].

To connect a 3D imaging probe with $32 \times 32 = 1024$ channels, assuming the same sampling frequency and resolution, the bandwidth requirements would increase by 16 times to 245 Gbps. Even with a more compact encoding, as allowed by an Aurora 64b66b module [17], the raw bandwidth would reach 256 Gbps and therefore would require a minimum of four QSFP+ connections. This is not presently implementable on the KCU105 development board due to the hard-wired connection of the FPGA transceivers to a variety of different interfaces, which poses a limit of two QSFP+ connections. Nonetheless, by developing a custom board, it is possible to use all 20 GTH transceivers of the KU040 FPGA towards the probe interface, resulting in up to five QSFP+ interfaces with an aggregate maximum bandwidth of 320 Gbps, which meets the requirements of 3D imaging. For even further scalability to higher channel counts or sampling rates, other FPGA chips of the Kintex Ultrascale family offer up to 64 GTH transceivers, thus providing at least three times more bandwidth.

2) *Offline Data Input by Ethernet Connection:* To be able to perform development, debugging and quality evaluation, our design also supports off-line data inputs, which could come from a reference simulation - with the help of Field-II simulator [18] - or a previously-acquired dataset. For ease of connection to the development system, e.g. a laptop, a Gigabit Ethernet port is supported, together with a communication protocol. At the current time, this interface is our only option to test 3D imaging, since we have no access to 3D imaging probes.

C. Probe Automatic Positioning by Gyroscopes and Accelerometers

Normally, the sonographer has full control over the positioning of the probe and there is a direct hand-eye coordination allowing the specialist to seek for anatomical features. Telesonography severs this connection and makes it time-consuming or problematic for a remote doctor to even identify the location of a scan. For example, the two kidneys would be essentially indistinguishable from each other if acquired in unlabeled scans. Therefore, it is imperative for the scans to be somehow position-tagged. Yet, a lightly trained operator may not have the necessary expertise or precision.

A possible solution is a probe positioning module, mechanically attached to the probe and capable of automatically detecting its own position. We choose an iNemo-M1 IMU from STMicroelectronics [19]. It is equipped with 3D gyroscope, 3D accelerometer, and 3D magnetometer in a compact package of $13 \times 13 \times 2\text{mm}$, which is tiny enough to be easily soldered into a transducer head. We have programmed the microcontroller of the iNemo-M1 with an Attitude and Heading Reference System (AHRS) algorithm developed in *C*. Since the acquired motion signals only convey acceleration and orientation, but not position, the AHRS algorithm is used to define the probe position $P_p(t)$ [20]. Eq. (1) shows the calculation of $P_p(t)$, where a_m is the measured acceleration, g is the gravitational acceleration (9.81m/s^2), and t_0 and t are the start and end time of measuring a_m , respectively. R_S^I represents the rotation matrix of the body frame of the sensor relative to the inertial frame¹. We have computed R_S^I based on an open source AHRS algorithm [21], [22].

$$P_p(t) = \int_{t_0}^t \int_{t_0}^t a_I(t') dt' = \int_{t_0}^t \int_{t_0}^t \left(R_S^I a_m(t') + \begin{pmatrix} 0 \\ 0 \\ g \end{pmatrix} \right) dt' \quad (1)$$

¹The frame is the origin and the three physical reference points that define the location and orientation of the coordinate system.

D. Volume Reconstruction

1) *Single-chip Beamforming*: BF is the kernel of any US image reconstruction. The received echoes are first of all weighted in two different ways: (i) Time Gain Compensation (TGC), which counterbalances the depth- and time-dependent attenuation due to the waves traveling through tissues, and (ii) apodization. Apodization compensates the antenna-array-like effects of the transducer matrix, which, if uncorrected, would result in side lobes of comparable amplitude to the main lobe, reducing lateral resolution [23].

The weighted echoes are then summed along different delay trajectories; for every focal point of the output image, a delay trajectory and a summation are needed. The delay trajectory is the two-way time of travel from the acoustic wave origin O , the focal point S , and back to each of the probe elements D . The calculation of the delay between S and each D represents the most difficult processing step, requiring 200 billion square root calculations per second with the parameters in Table I and assuming a reconstruction rate of 100 volumes/s.

Finally, as the BF image is still in the RF range, a demodulation step is needed to bring it to baseband frequency.

For each of these steps, we have devised optimized processing blocks [16], [24]–[26], suitable for low-power and FPGA implementation. In particular, we (i) pre-calculate a TGC table; (ii) perform apodization with a static table, which is also pre-calculated; (iii) approximate the delay calculation with a Taylor expansion, which introduces inaccuracies but simplifies the square root operations into additions between tables of pre-computed values; (iv) efficiently store the echo samples on a compact set of BRAMs, leveraging the dual-ported memory architecture of the Xilinx UltraScale family; (v) perform demodulation with a compact circuit consisting of an absolute value operation followed by 4th-order FIR low-pass filtering.

E. Post-processing and Visualization

1) *Cross-sectional Scan-Conversion and HDMI Support*: A beamformed US image is not interpretable by the human eye due to its extreme dynamic range. Log Compression (LC), i.e. a logarithmic brightness manipulation, needs to be applied to improve the image readability. Moreover, since most US imaging techniques - including ours - reconstruct images in polar coordinates, a transformation to Cartesian coordinates is needed. This transformation, which is called Scan Conversion (SC), allows the image to be shown on a grid of pixels on a screen. Full 3D SC is also computationally and materially expensive process to fit on FPGA. However, since our platform targets teleonography, 3D SC only needs to be performed at the hospital side, while the operator just needs to receive basic visual feedback, for which 2D cross-sections suffice. Therefore, we have implemented on our board a cross-sectional post-processing unit. The scan-converted image can then be displayed on a screen, which in the current prototype can be connected to an HDMI output. We have used the Analog Devices ADV7511 [27] module on our KCU105 board for HDMI transmission. The proposed post-processing block allows the operator, using the on-board push-buttons, to choose which cross-section of the volume to be scan-converted and displayed. Moreover, it gives the operator a full control to adjust the image brightness and compression degree of the LC through a developed GUI [16].

IV. EXPERIMENTAL RESULTS

A. FPGA Architectural Results

We have deployed the proposed design onto the Kintex UltraScale KCU105 FPGA board. Table II shows the resource utilization of the 1024-channel 3D imager on the KCU105 board. Since the operating frequency of the beamformer, including the demodulator, is 133 MHz and the reconstruction rate is 1 voxel/s, the theoretical reconstruction rate of a volume of $64 \times 64 \times 500$ voxels is 66 volumes/sec (vps). This theoretical rate is currently limited to 0.02 vps in practice, mostly due to the Ethernet bandwidth bottleneck for data communication, which is a necessary workaround due to the lack of availability of a native 3D imaging probe. The low refresh rate is partially also due

to inefficiencies in the SC block, the optimization of which is part of our on-going optimization work. We estimate that the system can be further upscaled to 90×90 channels on a larger, commercially available, FPGA like the Virtex XCVU190; this channel count is by far beyond current state of the art, but would also imply higher cost and power consumption.

B. Quality Assessment

We have assessed and quantified in detail the image reconstruction quality of our system in [26], [28]. The assessment includes Point Spread Function (PSF) evaluation, numerical inaccuracy quantification, and statistical inaccuracy distribution within the volume. The result of this study is that image quality is only marginally affected compared to a golden reference BF algorithm, and only in areas (the shallowest region and the edges) which are usually less critical for diagnosis.

For what concerns our probe positioning system, the reconstructed position's accuracy is a function of the acquisition time, since any measurement error is integrated over time. We also observed that the accuracy depends on the probe movement pattern; in particular, fast movement with high accelerations is more precisely detected than smooth, constant-speed gestures. Based on our initial results, for an acquisition time of about 10s, the system - calibrated at a reference starting position, e.g. the sternum - is sufficiently precise to detect basic positional coordinates, like "left" as opposed to "right" of the body, or "thorax" as opposed to "abdomen". This can reduce misinterpretation chances, however, more precise positioning will need to rely on higher-accuracy systems, e.g. visual or RF tracking. Alternatively, the scans can be manually tagged. Yet another possible approach is to design a companion mobile application which, for a variety of relevant diagnostic scans - e.g., "suspected abdominal bleeding" - visually instructs the operator on where to position the probe.

C. System Setup

Figure 2 shows the setup of the proposed teleonographic platform. The probe is connected to the KCU105 board of our imager using a Quad enhanced Small Form-factor Pluggable (QSFP+) port. Since the KCU105 does not natively support QSFP+, an FPGA Mezzanine Card (FMC) [29] is used as an interface board. The imager output is then sent over an HDMI transmitter to a screen. The platform is also able to communicate with a laptop using Ethernet, for simulated data processing. For demonstration purposes, we show in Fig. 2 both the laptop and the screen as a single tablet-like device, which would be the future setup.

D. Platform Portability and Estimated Cost

Our current prototype leverages a standard Xilinx KCU105 development board, which is obviously not optimized for price or size. Nonetheless, we note that we could acquire the board for about 3000€. The QSFP+ Mezzanine Card (FPGA) interface needed to add a high-speed optical interface retails for about 500€, while the iNemo-M1 IMU to track the probe position can be acquired for less than 100€. Even including a screen, a battery and casing, the total bill of materials is therefore less than 4000€, excluding the probe. We estimate that a custom, optimized, mass-produced PCB would probably reduce the design cost in half. This compares extremely

TABLE II
IMAGER RESOURCE UTILIZATION.
*Kintex UltraScale KU040 implementation results.
**Virtex UltraScale XCVU190 extrapolated results.

Supported Channels	Logic LUTs	Regs	BRAM	DSP	Clock	Theo. Rate
$32 \times 32^*$	85%	34.5%	97.1%	7.2%	133 MHz	66 vps
$90 \times 90^{**}$	84.3%	31.5%	89.4%	7.7%	133 MHz	66 vps




Fig. 2. Setup of the teleosonographic-capable 3D US imager. *The probe on the right is a 1D probe [14] for proof-of-concept and to be replaced with a matrix probe whenever available.

favorably to commercial 3D systems, which retail for up to 100'000€, and potentially enables a wide adoption.

At the moment, the size of the prototype is dominated by that of the KCU105 board, i.e. $26.7\text{cm} \times 14\text{cm} \times 0.16\text{cm}$. Even without any redesign, this is clearly compatible with a portable device. The estimated power consumption is comparable to that of a laptop CPU, suggesting that full-day operation with an inexpensive and compact battery is possible. The iNemo-M1 would need to be integrated into the transducer head, and was in fact chosen partially due to its small size ($13 \times 13 \times 2\text{mm}$), which makes the integration easily possible.

Overall, we believe that our prototype, in addition to being the first attempt at demonstrating a feasible teleosonography platform and supporting an extremely high channel count, achieves an unprecedented performance/power/cost trade-off.

V. CONCLUSION AND FUTURE WORK

We have proposed the first teleosonography-capable US platform that supports up to 1024 channels on a single FPGA, with a material cost of less than 4000€, and an estimated power consumption of 6.1W. Our design is a complete 3D US imager including signal acquisition, 3D reconstruction, and post-processing. We support real-time data input from a probe over a high-bandwidth optical connection, as well as a development and debug Ethernet interface. The platform can output US images on an HDMI screen. Further, we have provided an automatic location detection facility for the probe, using an IMU. This further facilitates the decoupling of acquisition and diagnosis, enabling teleosonography and the related potential societal benefits. We plan on iterating on the current prototype, optimizing its performance, resource utilization, power consumption, and features.

ACKNOWLEDGMENT

The authors acknowledge Swiss Confederation funding through the UltrasoundToGo project of the Nano-Tera.ch initiative.

REFERENCES

- [1] Xilinx Inc., “Xilinx kintex ultraScale FPGA KCU105 evaluation kit,” 2017, <https://www.xilinx.com/products/boards-and-kits/kcu105.html>.
- [2] Xilinx Inc., “Ultrascale FPGA: Product tables and product selection guide,” 2016, <http://www.xilinx.com/support/documentation/selection-guides/ultrascale-fpga-product-selection-guide.pdf#KU>.
- [3] J. Larson, “2-d phased array ultrasound imaging system with distributed phasing,” Jul. 20 1993, uS Patent 5,229,933. [Online]. Available: <https://www.google.ch/patents/US5229933>
- [4] “EPIQ 7 ultrasound system,” Philips Healthcare, 2017. [Online]. Available: <http://www.medical.philips.com/main/products/ultrasound/systems/epiq7/>
- [5] Philips Healthcare, “Cx50 xMATRIX,” 2017, <http://www.usa.philips.com/healthcare/product/HC795076CA/cx50-xmatrix-cardiology-ultrasound-machine>.
- [6] Cephasonics, “cQuest Griffin - Product Brief,” Catalog, 2015, http://pdf.medicaexpo.com/pdf/cephasonics/cquest-griffin/92255-171445-_2.html.
- [7] R. J. Tobias and M. E. Schafer, “System development improvements using stackable high-channel-count ultrasound hardware,” in *2014 IEEE International Ultrasonics Symposium (IUS 2014)*, Sept 2014.
- [8] J. Jensen, H. Holtén-Lund, R. Nilsson, M. Hansen, U. Larsen, R. Domstén, B. Tomov, M. Stuart, S. Nikolov, M. Pihl, Y. Du, J. Rasmussen, and M. Rasmussen, “Sarus: A synthetic aperture real-time ultrasound system,” *Ultrasonics, Ferroelectrics, and Frequency Control, IEEE Transactions on*, vol. 60, no. 9, pp. 1838–1852, Sep 2013.
- [9] R. Sampson, M. Yang, S. Wei, R. Jintamethasawat, B. Fowlkes, O. Kripfgans, C. Chakrabarti, and T. Wenisch, “FPGA implementation of low-power 3D ultrasound beamformer,” in *2015 IEEE International Ultrasonics Symposium (IUS 2015)*, Nov 2015.
- [10] E. Boni, L. Bassi, A. Dallai, F. Guidi, V. Meacci, A. Ramalli, S. Ricci, and P. Tortoli, “ULA-OP 256: A 256-channel open scanner for development and real-time implementation of new ultrasound methods,” *Ultrasonics, Ferroelectrics, and Frequency Control, IEEE Transactions on*, in press 2016.
- [11] GE Healthcare, “Voluson i,” 2017, http://www3.gehealthcare.co.uk/en-gb/products/categories/ultrasound/voluson/voluson_i.
- [12] kpi Healthcare, “Samsung MySono U6 ultrasound machine,” 2017, <http://www.kpiultrasound.com/medical-equipment/ultrasound-scanners/samsung-ultrasound/samsung-mysono-u6-portable/>.
- [13] Chison Medical Imaging Co., Ltd., “Chison q9,” 2014, http://www.chison.com/products_detail.asp?id=168.
- [14] 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 *Proceedings of the 2017 IEEE International Symposium on Circuits and Systems (ISCAS 2017)*, 2017.
- [15] Xilinx Inc., “Aurora 8B/10B,” 2017, <https://www.xilinx.com/products/intellectual-property/aurora8b10b.html>.
- [16] A. Ibrahim, W. Simon, D. Doy, E. Pignat, F. Angiolini, M. Arditì, J.-P. Thiran, and G. De Micheli, “Single-FPGA complete 3D and 2D medical ultrasound imager,” in *Proceedings of the Conference on Design and Architectures for Signal and Image Processing (DASIP 2017)*, 2017.
- [17] Xilinx Inc., “Aurora 64B/66B,” 2017, <https://www.xilinx.com/products/intellectual-property/aurora64b66b.html>.
- [18] J. Jensen, “Field: A program for simulating ultrasound systems,” in *10th Nordic-Baltic Conference on Biomedical Imaging Published in Medical and Biological Engineering and Computing*, vol. 34, Supplement 1, 1996, pp. 351–353.
- [19] STMicroelectronics Inc., “iNemo-M1 System-On-Board,” 2017, <https://www.digikey.com/en/product-highlight/s/stmicroelectronics/inemo-m1-system-on-board>.
- [20] CHRbotics LLC., “Using accelerometers to estimate position and velocity,” <http://www.chrobotics.com/library/accel-position-velocity>.
- [21] x-io Technologies, “Open source IMU and AHRS algorithms,” 2012, <http://x-io.co.uk/open-source-imu-and-ahrs-algorithms/>.
- [22] S. Madgwick, A. Harrison, and R. Vaidyanathan, “Estimation of IMU and MARG orientation using a gradient descent algorithm,” in *Proceedings of the 2011 IEEE International Conference on Rehabilitation Robotics (ICORR)*, 2011.
- [23] C. Daft and W. Engeler, “Windowing of wide-band ultrasound transducers,” in *Ultrasonics Symposium, 1996. Proceedings., 1996 IEEE*, vol. 2, IEEE, 1996, pp. 1541–1544.
- [24] A. Ibrahim, P. A. Hager, A. Bartolini, F. Angiolini, M. Arditì, L. Benini, and G. De Micheli, “Tackling the bottleneck of delay tables in 3D ultrasound imaging,” in *Proceedings of the 2015 Design Automation and Test in Europe (DATE 2015) Conference*, March 2015, pp. 1683 – 1688.
- [25] F. Angiolini, A. Ibrahim, W. Simon, A. C. Yüzügüler, M. Arditì, J.-P. Thiran, and G. De Micheli, “1024-Channel 3D ultrasound digital beamformer in a single 5W FPGA,” in *Proceedings of the 2017 Design Automation and Test in Europe (DATE 2017) Conference*, March 2017.
- [26] A. Ibrahim, P. A. Hager, A. Bartolini, F. Angiolini, M. Arditì, J.-P. Thiran, L. Benini, and G. De Micheli, “Efficient sample delay calculation for 2d and 3d ultrasound imaging,” to appear in *IEEE Transactions on Biomedical Circuits and Systems (TCAS)*, 2017.
- [27] Analog Devices Inc., “ADV7511: 225 mhz, high performance hdmi transmitter with arc,” 2017, <http://www.analog.com/en/products/audio-video/analoghdmidvi-interfaces/hdmidvi-transmitters/adv7511.html#product-overview>.
- [28] A. Ibrahim, F. Angiolini, M. Arditì, J.-P. Thiran, and G. De Micheli, “Apodization scheme for hardware-efficient beamformer,” in *Proceedings of the 12th Conference on PhD Research in Microelectronics and Electronics (PRIME)*, 2016.
- [29] HiTech Global LLC, “2-port QSFP+ (40G or 56G) FMC module (vita57.1),” http://www.hitechglobal.com/FMCModules/FMC_2QSFP+.htm.