

This is the peer reviewed version of the following article:

ECHOES: a 200 GOPS/W Frequency Domain SoC with FFT Processor and I2S DSP for Flexible Data Acquisition from Microphone Arrays / Sinigaglia, M.; Bertaccini, L.; Valente, L.; Garofalo, A.; Benatti, S.; Benini, L.; Conti, F.; Rossi, D.. - 2023-(2023), pp. 1-5. (56th IEEE International Symposium on Circuits and Systems, ISCAS 2023 usa 2023) [10.1109/ISCAS46773.2023.10181862].

Institute of Electrical and Electronics Engineers Inc.

Terms of use:

The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. For all terms of use and more information see the publisher's website.

01/05/2026 00:52

(Article begins on next page)

ECHOES: a 200 GOPS/W Frequency Domain SoC with FFT Processor and I²S DSP for Flexible Data Acquisition from Microphone Arrays

Mattia Sinigaglia*, Luca Bertaccini[†], Luca Valente*, Angelo Garofalo*[†], Simone Benatti[‡], Luca Benini*[†],
Francesco Conti*, and Davide Rossi*

*Department of Electrical, Electronic and Information Engineering (DEI), University of Bologna, Italy

[†]Integrated Systems Laboratory (IIS), ETH Zürich, Switzerland

[‡]Department of Science and Methods for Engineering (DISMI), University of Modena e Reggio Emilia, Italy

Abstract—Emerging applications in the IoT domain require ultra-low-power and high-performance end-nodes to deal with complex near-sensor-data analytics. Domains such as audio, radar, and Structural Health Monitoring require many computations to be performed in the frequency domain rather than in the time domain. We present ECHOES, a System-On-a-Chip (SoC) composed of a RISC-V core enhanced with fixed- and floating-point digital signal processing (DSP) extensions and a Fast-Fourier Transform (FFT) hardware accelerator targeting emerging frequency-domain application. The proposed SoC features an autonomous I/O engine supporting a wide set of peripherals, including Ultra-Low-Power radars, MEMS, and digital microphones over I²S protocol with full-duplex Time Division Multiplexing DSP mode, making ECHOES the first open-source SoC which offers this functionality enabling simultaneous communication with up to 16 I/Os devices. ECHOES, fabricated with 65nm CMOS technology, reaches a peak performance of 0.16 GFLOPS and a peak energy efficiency of 9.68 GFLOPS/W on a wide range of floating and fixed-point general-purpose DSP kernels. The FFT accelerator achieves performance up to 10.16 GOPS with an efficiency of 199.8 GOPS/W, improving performance and efficiency by up to 41.1× and 11.2×, respectively, over its software implementation of this critical task for frequency domain processing.

Index Terms—Frequency Domain SoC, FFT processor, I²S, TDM, microphone array

I. INTRODUCTION AND RELATED WORK

A wide range of applications in the IoT domain requires powerful end nodes capable of dealing with the increasing complexity of near-sensor-data analytics algorithms. A common feature of these battery-powered systems is the need to operate within a power envelope of less than 100 mW, which, coupled with the high-performance requirements of the algorithms, leads to extremely tight energy efficiency requirements. Several emerging application fields such as audio, gesture recognition, or Structural Health Monitoring (SHM) perform computations such as noise canceling, filtering, and equalization in the frequency domain rather than in the time domain, exploiting emerging sensors such as MEMS, radars, and microphone arrays.

Due to the nature of these applications, devices must be powerful but, at the same time, must also have a tiny physical footprint and power consumption. An extreme example is that of hearing aids [1]. The minimum power for a given function can be achieved with fully specialized dedicated hardware [2], but this approach comes with limits in flexibility, and post-fabrication updates [3]. An emerging trend is to couple flexibility with efficiency, fully programmable microcontroller cores enhanced with Digital Signal Processing (DSP)-oriented

instruction set extensions. This solution provides enough flexibility to deal with complex and fast-evolving frequency-domain algorithms. In this way, the same device can target various application fields. For example, structural monitoring analyzing vibrations [4], noise cancellation for consumer applications such as headphones [5] [6] or smart glasses [7], gesture recognition, as well as movement and object detection based on radar processing [8] [9]. Since these devices must be cheap and produced at a very large scale to be economically viable, this "one-size-fits-many" design approach maximizes the usefulness of emerging DSP architectures.

On the other hand, to meet strict application performance requirements, DSP devices must run fast and at low power selected key kernels that are 1) highly computationally intensive and 2) used repeatedly. A prime example is that of Fast-Fourier Transform (FFT), which plays a key role in DSP algorithms, often taking a large portion of the overall computation time. Augmenting a flexible architecture with dedicated FFT accelerators is a good strategy to increase performance and efficiency. Most FFT accelerators [10] [11] have been designed with internal memories to host and internally reshuffle butterfly data. However, such internal buffers usually occupy most of the area dedicated to the accelerator, increasing their cost [12]. To mitigate this issue, Baas [13] proposed an FFT accelerator including a small cache, while other implementations share the memory with cores such as [14]. Furthermore, when addressing a wide range of applications, FFT accelerators that support multiple data types are desirable. Using such designs gives a chance to lower the precision, whenever possible, to improve performance and energy efficiency.

Among audio applications, microphone arrays are becoming widespread as many applications require the exploration and characterization of background sound and/or multiple, spatially distributed noise sources. One way to overcome the limitations of perception of speech intelligibility in time and space in varying noisy environments is the use of microphone arrays, which provides many advantages: directional reception of the sound, spatial localization of the target speaker, or noise suppression of point sources [15]. Beamforming algorithms take advantage of multi-microphone arrays to perform suppression of unwanted contamination sources [16]. Microphones in digital devices use Inter-Integrated Circuit (IC) Sound I²S protocol [17] [18], which was originally developed for transmitting stereo audio data through a serial line. To perform multi-microphone acquisition, many System-On-a-Chip (SoC)s provide multiple I²S peripherals. This technique leads to high-area occupation and can not be used in small devices, which are often pad-limited. Thus, SoCs with strict




Fig. 1: a) Overview of ECHOES SoC and details of: b) FFT HWPE accelerator c) I²S with DSP implementation

power and area constraints have to deal with this problem differently. A popular technique is Time Division Multiplexing (TDM). TDM allows multiple devices to share the same communication line; each device uses the line for a specific time fraction, and the protocol has to ensure there are no collisions. The limitation of this technique applied to the I²S protocol is the latency introduced by the multiplexing of multiple sources. Many devices follow the DSP mode implementation developed by NXP for implementing low-latency TDM [19].

Some solutions based on RISC-V targeting the audio domain recently appeared on the market. Notable examples are GreenWaves Technologies GAP9 [20] and Telink TLR9 [21]. Unfortunately, it is impossible to compare with these commercial SoCs since their internal architecture and techniques used to achieve high efficiency are not openly described. From the academic viewpoint, the most recent architectures are those proposed by Nathaniel et al. [22] implementing flexible acceleration through coarse grain reconfigurable arrays (CGRA), and by Jun et al. hosting a hardware accelerator for parallel convolution operation [23] targeting Active Noise Control ANC algorithms. On the other hand, while the proposed architectures are specifically designed to target the audio domain, the proposed SoC specifically targets frequency domain applications by coupling a dedicated FFT processor with a general-purpose DSP for the flexible acceleration of frequency domain functions.

We present ECHOES, a 65 nm SoC specialized for frequency-domain DSP, centered around an industrially verified extended-ISA RISC-V CV32 core for fixed- and floating-point energy efficient DSP. ECHOES introduces two main points of novelty. First, it couples the CV32 core with a shared-memory, flexible data-width FFT hardware accelerator to boost conversions to/from the frequency domain, which is a predominant bottleneck in DSP algorithms. Second, it is the first open-source SoC implementing the high configurable full-duplex I²S TDM DSP mode for communication with up to 16 digital I/Os audio devices. ECHOES can be used as the central core for DSP applications or as an additional device offering FFT hardware capabilities. Benchmarked a wide range of floating and fixed point relevant kernels for audio processing, ECHOES achieves a peak performance of 0.18 GFLOPS and a peak efficiency of 9.68 GFLOPS/W when using the CV32 core. The FFT accelerator for frequency-domain processing achieves performance up to 10.16 GOPS with an efficiency of 199.8 GOPS/W which is 11.2× higher compared to the software implementation and 41.1× faster. The hardware and software described in this work are open-source, intending to support and boost an innovation ecosystem focusing on ULP computing for the IoT landscape.

II. SOC OVERVIEW

Fig. 1 a) shows the architecture of ECHOES. The SoC consists of an in-order CV32E40P [24] 32-bit RISC-V core with

a 4-stage pipeline that implements the RV32IMFC instruction set architecture. It features extensions for both floating- and fixed-point energy-efficient DSP such as hardware loops, post-increment LD/ST, and Single Instruction Multiple Data (SIMD) such as dot products operating on narrow 16b and 8b data types [25] [26]. The core is extended with a Floating Point Unit (FPU) supporting single-precision floating-point FP32 arithmetic, addition, subtraction, square root, and division, as well as Fused Multiply and Accumulate (FMA), which is a critical kernel for frequency-domain digital signal processing.

The SoC is enhanced with an autonomous I/O subsystem coupling an I/O DMA tightly-coupled with a multi-banked system memory [27]. The memory hierarchy of the system is composed of a 256 kB L2 memory used to share data among the core, the peripherals, and the accelerator. The L2 memory is organized in 16 word-level interleaved banks to deliver up to 22.4 GB/S to the master resources. The I/O engine supports a broad set of peripherals such as QSPI, SPI, UART, and I²C that can acquire data from multiple sensors. It also includes an 800 Mb/s DDR interface supporting external IoT DRAMs such as Cypress Semiconductor’s HyperRAM. Two configurable FLLs provide a dedicated clock source to the processor and peripherals.

The following sections introduce the two architectural contributions of this work relevant for frequency-domain processing and audio applications: the tightly-coupled *FFT Accelerator* and the *I²S DSP Time Division Multiplexing*.

A. FFT Accelerator

To enhance the DSP capabilities of ECHOES, we integrate a radix-2 decimation-in-time FFT Hardware Processing Engine (HWPE) into the SoC. HWPEs are a class of hardware accelerators that share the memory with one or multiple compute cores, thus enabling efficient cooperation between the general-purpose and the domain-specific part of the architecture. Such an approach is uncommon in traditional FFT accelerators, which are often decoupled from the core, requiring data copies and large internal buffers.

Since the butterfly organization of the FFT algorithm easily leads to systematic banking conflicts, the HWPE implements the scheme recently proposed by [12] to reorder the butterfly sequence and the memory accesses. This scheme requires access to at least 16 distinct memory banks. It allows the accelerator to load and store data at consecutive memory locations by moving sets of right or left butterfly wings. The samples are then reordered before and after the butterfly computation using internal registers. In this way, banking conflicts can only appear during the final bit-reversed reordering required by the FFT algorithm; they are handled by the accelerator pipeline for a single cycle.

The FFT accelerator works with fixed-point complex data whose real and imaginary parts are represented with 32/16/8 bits. We will refer to these data types as C64/C32/C16. The




Fig. 2: Example of I²S Time Division Multiplexing with 4 devices: a) I²S TDM b) I²S TDM DSP Mode. Each color represents a device. The frame length is n , while the channel length is $k = n/2$.

maximum number of FFT points supported is 512/1024/2048, respectively for $C64/C32/C16$. The accelerator is programmed by the core, responsible for providing the number of FFT points, the data type, the memory address at which the input samples are stored, and the start command. The HWPE contains eight 32-bit memory ports, of which four are used as input and four as output ports, and it is organized in submodules as shown in Fig. 1 c).

The Butterfly Unit is composed of one $C64$, one $C32$, and two $C16$ butterfly engines. As wider butterfly engines can be reused for lower-precision computations, the Butterfly Unit can compute up to $1/2/4$ $C64/C32/C16$ butterflies per cycle. Two sets of four $C64$ Butterfly Registers are located around the Butterfly Unit to reorder inputs and outputs, while the twiddle factors are stored in a lookup table.

B. I²S DSP Time Division Multiplexing

I²S is a standard for digital audio acquisition that defines three main signals: BCLK represents the clock signal; SD (Serial Data) transports the audio data; FSYNC (Frame Synch) selects the stereo channel of transmitted data (Left or Right). The peripheral can act both as a master and a slave. The I²S protocol specifies that the master serially drives data over the SD lines on the falling edges of the BCLK, and the slave samples them on the rising edge. The FSYNC signal discriminates between the Left and Right channels.

In the TDM I²S implementation, the multiplexed devices completely bind the time needed to receive both channels from a specific source. Each device of the array first sends the L channel and then the R channel, following the connection order of the array. This order implies that the number of devices introduces significant latency in communication – a very undesirable effect, particularly for audio processing, which is highly latency sensitive. Fig. 2 a) shows an example of the I²S TDM implementation with 4 devices. Considering that each frame is composed of $nBits$ and the number of the multiplexed devices is K , each frame of a source has a latency of $\frac{nBits}{2} \times (K + 1) \times Tclk$.

To reduce audio signal acquisition latency, ECHOES implements a more specialized I²S protocol dealing with the parallel acquisition from multiple microphones: TDM DSP mode. Devices supporting DSP mode transmit L/R channels one after the other immediately when the FSYNC is asserted. Fig. 2 b) shows the implementation of the DSP TDM mode with an example of 4 devices. In this implementation, each device occupies a specific slot during which it sends the entire frame, and the latency to complete the transmission is reduced to a fixed $nBits \times Tclk$.

Fig. 1 c) shows the block diagram of the I²S peripheral. It is composed of two interfaces for transmitting and receiving audio data. Each interface can be programmed to act as DSP TDM or standard I²S. Both have dedicated I²S buses, which allow for independent and full-duplex communications. The core programs the memory-mapped 32-bit register file,

allowing each interface’s independent configuration flexibility. By scaling the peripheral clock, the peripheral features 2 clock dividers for generating the interfaces’ BCLK signals. The four dedicated and independent FSYNC blocks generate the FSYNC signal accordingly to the I²S and DSP modes for each interface.

This operation mode is designed for multiplexing 16 devices that work with a maximum 48 kHz sample rate. The BCLK frequency of the TDM DSP mode depends on the combination of $nDevices \times frameBits \times sampleRate$. A SoC frequency of around 25 MHz (much lower than the maximum achievable in ECHOES) is sufficient to support the best available rate when samples have a frame length of 32 bits. Thanks to the high frequency reached by ECHOES, it is possible to read and compute data before $nDevices$ new data have been received and stored. To deal with different implementations of the DSP mode, ECHOES can also be programmed to choose the clock polarity and program the frame alignment.

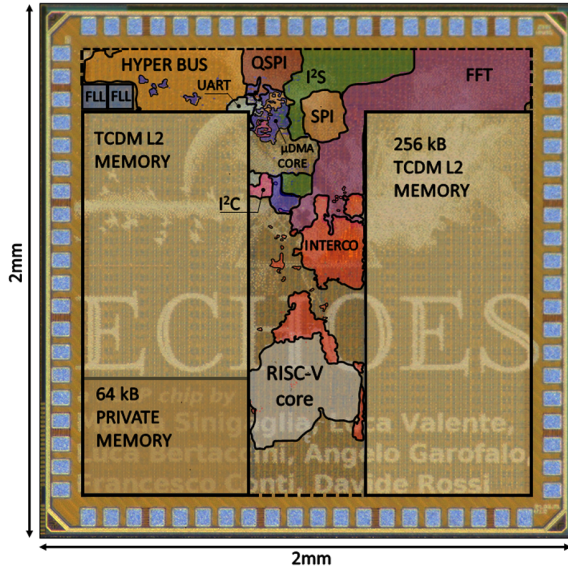
III. PHYSICAL IMPLEMENTATION AND MEASUREMENTS

ECHOES has been implemented in TSMC 65 nm technology performing synthesis with Synopsys Design Compiler [28], Place & Route with Cadence Innovus [29], and verification and sign-off with Calibre [30]. The die area is 4 mm², partitioned as shown in Fig. 3 where the majority of the area is occupied by 256 kB L2 Memory and the 64 kB private memory. The FFT accelerator occupies 6%, the I²S occupies 2% of the total area, and the core occupies 3%.

The ECHOES chip has been tested and characterized using an Advantest SoC hp9300 integrated circuit testing device. The measurements are taken by executing Matmul and both software and hardware-accelerated FFTs on the chip. Fig. 4 reports the maximum operating frequency and the power consumption over 0.9 V to 1.2 V for representative of many compute-intensive workloads in the field of audio processing. The Matmuls kernels have been performed using 16×16 matrixes and consume respectively 33 mW at 1.2 V for INT32, and 33.6 mW at 1.2 V for FP32. The power consumption of the FFT accelerator has been measured executing $C64/C32/C16$ workloads respectively on 512/1024/2048 FFT points. The peak power consumption of the FFT reaches 133.5 mW at 1.2 V when executing the FFT $C64$ on 512pt.

IV. BENCHMARKING

To demonstrate the capabilities and flexibility of ECHOES, we used benchmark representative of typical frequency-domain and audio DSP workloads. Fig. 5 shows the performance and energy efficiency of the RISC-V core computing FP32 DSP kernels. Thanks to the high versatility of the programmable core, ECHOES can execute a variety of workloads from frequency- and time-domain processing applications. Fig. 6 reports the performance and the energy efficiency of FFT kernels computed using the accelerator and the core, as well as the related speed-up. Thanks to the multiple data sizes



Technology	Chip Area	SRAM
CMOS 65nm	4mm ²	256 kB
Vdd Range	Frequency Range	Power Envelope
0.9 - 1.2V	150 - 350MHz	133.5 mW

Fig. 3: Chip micrograph and specifications




Fig. 4: Voltage sweep vs. max frequency vs. power consumption.

supported by the FFT accelerator, the precision of a specific kernel can be tuned to the application requirements, enhancing the performance and energy efficiency by 2× when moving from C64 to C32 formats.

To put our results in perspective, we benchmark the Mel-Frequency Cepstral Coefficients (MFCC) feature extractor and compare it against the solution proposed in [31]. Contrarily to [31], where the execution is offloaded to a parallel 8 cores cluster, in ECHOES, we exploit the tightly-coupled cooperation between the single general-purpose core and the specialized FFT accelerator, leading to a much more compact form factor of the system (4mm² vs 10mm²). Despite the fact that the inference of a single MFCC on ECHOES is executed in 120μs, 1.3× slower than [31], due to non-FFT kernels executed on a single core instead of a multi-core cluster, our solution outperforms [31] by 2× in area efficiency, despite the less scaled technology node used for implementation.

V. CONCLUSION

We presented ECHOES, a SoC composed of a RISC-V core enhanced with fixed- and floating-point DSP exten-


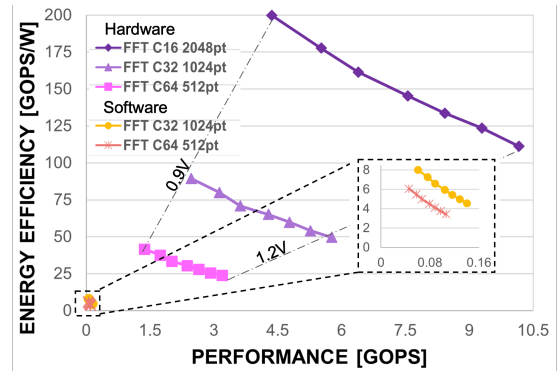


Fig. 5: Performance and energy efficiency of DSP kernels on ECHOES.



Kernel	Peak Efficiency SW	Peak Efficiency HW	Gain
FFT C32 1024pt	8 GOPS/W	89.5 GOPS/W	11.2 ×
FFT C64 512pt	6.1 GOPS/W	41.6 GOPS/W	6.9 ×

Kernel	Peak Performance SW	Peak Performance HW	Speedup
FFT C32 1024pt	0.1 GOPS	5.8 GOPS	41.1 ×
FFT C64 512pt	0.1 GOPS	3.2 GOPS	30.1 ×

Fig. 6: Performance and Energy Efficiency of FFT kernels on ECHOES. Speedup and energy gain of FFT computation on the hardware accelerator.

sions and an FFT hardware accelerator targeting emerging frequency-domain applications. The proposed SoC features an autonomous I/O engine able to interface to a wide range of sensors such as MEMS, ULP radars, and digital microphones over I²S protocol with full-duplex TDM DSP enabling simultaneous communication with up to 16 I/Os devices. Thanks to the many configurations implemented, ECHOES reaches extreme flexibility and makes it the first open-source¹ SoC, which implements all the TDM DSP configurations available in the market. The proposed SoC, fabricated in TSMC 65 nm technology, can achieve FP32 peak performance up to 0.18 GFLOPS at 1.2 V with an efficiency of 9.68 GFLOPS/W at 0.9 V. The FFT accelerator for frequency domain DSP processing achieves performance up to 10.16 GOPS at 1.2 V with an efficiency of 199.8 GOPS/W at 0.9 V which is 11.2× higher compared to the software implementation and 41.1× faster.

ACKNOWLEDGEMENT

Supported in part through funding from the ECSEL Joint Undertaking (JU) under grant agreement No 877056 (FRAC-TAL). The JU receives support from the European Union's Horizon 2020 research and innovation programme and Spain, Italy, Austria, Germany, Finland, Switzerland.

¹<https://github.com/pulp-platform>

REFERENCES

- [1] K. Signature. (2022) Kirkland signature 10.0 with tcoil. [Online]. Available: {<https://ks10userportal.com/>}
- [2] L. Gerlach, G. Payá-Vayá, and H. Blume, “A survey on application specific processor architectures for digital hearing aids,” *Journal of Signal Processing Systems*, Mar 2021. [Online]. Available: <https://doi.org/10.1007/s11265-021-01648-0>
- [3] H. Puder, “Hearing aids: an overview of the state-of-the-art, challenges, and future trends of an interesting audio signal processing application,” in *2009 Proceedings of 6th International Symposium on Image and Signal Processing and Analysis*, 2009, pp. 1–6.
- [4] enDAQ. (2022) Shock, vibration and environmental sensors. [Online]. Available: {<https://endaq.com/>}
- [5] SONY. (2022) Wh-1000xm5 wireless industry leading noise canceling headphones. [Online]. Available: <https://electronics.sony.com/audio/headphones/all-headphones/p/wh1000xm5-b>
- [6] BOSE. (2022) Bose noise cancelling headphones 700 uc. [Online]. Available: \url{https://www.bose.com/en_us/index.html}
- [7] RAY-BAN. (2022) Ray-ban stories smart glasses. [Online]. Available: {<https://www.ray-ban.com/italy/ray-ban-stories>}
- [8] M. Scherer, M. Magno, J. Erb, P. Mayer, M. Eggimann, and L. Benini, “Tinyradarmn: Combining spatial and temporal convolutional neural networks for embedded gesture recognition with short range radars,” *IEEE Internet of Things Journal*, vol. 8, no. 13, pp. 10336–10346, 2021.
- [9] Y. Sun, T. Fei, X. Li, A. Warnecke, E. Warsitz, and N. Pohl, “Real-time radar-based gesture detection and recognition built in an edge-computing platform,” *IEEE Sensors Journal*, vol. 20, no. 18, pp. 10706–10716, 2020.
- [10] A. Wang, B. Richards, P. Dabbelt, H. Mao, S. Bailey, J. Han, E. Chang, J. Dunn, E. Alon, and B. Nikolić, “A 0.37mm² lte/wi-fi compatible, memory-based, runtime-reconfigurable 2n3m5k fft accelerator integrated with a risc-v core in 16nm finfet,” in *2017 IEEE Asian Solid-State Circuits Conference (A-SSCC)*, 2017, pp. 305–308.
- [11] L. Guo, Y. Tang, Y. Lei, Y. Dou, and J. Zhou, “Transpose-free variable-size fft accelerator based on-chip sram,” *IEICE Electronics Express*, vol. 11, no. 15, pp. 20140171–20140171, 2014.
- [12] L. Bertaccini, L. Benini, and F. Conti, “To buffer, or not to buffer? a case study on fft accelerators for ultra-low-power multicore clusters,” in *2021 IEEE 32nd International Conference on Application-Specific Systems, Architectures and Processors (ASAP)*. IEEE, 2021, pp. 1–8.
- [13] B. Baas, “A low-power, high-performance, 1024-point fft processor,” *IEEE Journal of Solid-State Circuits*, vol. 34, no. 3, pp. 380–387, 1999.
- [14] M. McKeown, “FFT implementation on the TMS320VC5505, TMS320C5505, and TMS320C5515 DSPs,” *Texas Instruments Incorporated, White Paper SPRABB6B*, 2010.
- [15] M. Stolbov, “Application of microphone arrays for distant speech capture,” *Scientific and Technical Journal of Information Technologies, Mechanics and Optics*, vol. 15, pp. 661–675, 06 2015.
- [16] A. Hidri, M. Souad, A. Alaqeeli, and H. Amiri, “Beamforming techniques for multichannel audio signal separation,” *International Journal of Digital Content Technology and its Applications*, vol. 6, 12 2012.
- [17] NXP. (2022) I2s bus specification. [Online]. Available: {<https://www.nxp.com/docs/en/user-manual/UM11732.pdf>}
- [18] A. DEVICES. (2012) Common inter-ic digital interfaces for audio data transfer. [Online]. Available: {<https://www.analog.com/media/en/technical-documentation/technical-articles/MS-2275.pdf>}
- [19] NXP. (2011) Freescale semiconductor application note AN3664: SGTL5000 I2S DSP Mode. [Online]. Available: <https://www.nxp.com/docs/en/application-note/AN3664.pdf>
- [20] G. Technologies. (2022) Gap9 processor for hearables and smart sensors. [Online]. Available: {https://greenwaves-technologies.com/processor_noise_cancellation/}
- [21] Telink. (2020) Introducing telink’s new risc-v-based tlr9 series. [Online]. Available: {<https://www.telink-semi.com/news/telink-risc-v-tlr9/>}
- [22] N. Bleier, M. Mubarik, S. Chakraborty, S. Kishore, and R. Kumar, “Rethinking programmable earable processors,” 06 2022, pp. 454–467.
- [23] J. Yuan, Q. Zhao, W. Wang, X. Meng, J. Li, and Q. Li, “Audio denoising coprocessor based on risc-v custom instruction set extension,” *Acoustics*, vol. 4, no. 3, pp. 538–553, 2022. [Online]. Available: <https://www.mdpi.com/2624-599X/4/3/33>
- [24] OpenHW. Open source hardware the new reality cv32e40p project. [Online]. Available: {<https://www.openhwgroup.org/resources/blog/open-source-hardware-the-new-reality-cv32e40p-project/>}
- [25] M. Gautschi, P. D. Schiavone, A. Traber, I. Loi, A. Pullini, D. Rossi, E. Flaman, F. K. Gürkaynak, and L. Benini, “Near-threshold risc-v core with dsp extensions for scalable iot endpoint devices,” *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 25, no. 10, pp. 2700–2713, 2017.
- [26] P. Davide Schiavone, F. Conti, D. Rossi, M. Gautschi, A. Pullini, E. Flaman, and L. Benini, “Slow and steady wins the race? a comparison of ultra-low-power risc-v cores for internet-of-things applications,” in *2017 27th International Symposium on Power and Timing Modeling, Optimization and Simulation (PATMOS)*, 2017, pp. 1–8.
- [27] A. Pullini, D. Rossi, G. Haugou, and L. Benini, “udma: An autonomous i/o subsystem for iot end-nodes,” in *2017 27th International Symposium on Power and Timing Modeling, Optimization and Simulation (PATMOS)*, 2017, pp. 1–8.
- [28] Synopsys. (2022) Synopsys — eda tools, semiconductor ip and application security solutions. [Online]. Available: <https://www.synopsys.com/>
- [29] Cadence. (2022) Innovus implementation system. [Online]. Available: https://www.cadence.com/ko_KR/home.html
- [30] Siemens. (2022) Calibre design solutions. [Online]. Available: https://www.cadence.com/ko_KR/home.html
- [31] M. Fariselli, M. Rusci, J. Cambonie, and E. Flaman, “Integer-only approximated mfcc for ultra-low power audio nn processing on multi-core mcus,” in *2021 IEEE 3rd International Conference on Artificial Intelligence Circuits and Systems (AICAS)*, 2021, pp. 1–4.