Aalborg Universitet

AALBORG
UNIVERSITY

Error Reduced Carry Prediction in Approximate Addition for Low-Frequency Sound Zones

Koch, Peter; Østergaard, Jan

Published in:
2022 18th Biennial Baltic Electronics Conference, BEC 2022

DOI (link to publication from Publisher):
10.1109/BEC56180.2022.9935613

Publication date:
2022

Document Version
Accepted author manuscript, peer reviewed version

Link to publication from Aalborg University

Citation for published version (APA):
Koch, P., & Østergaard, J. (2022). Error Reduced Carry Prediction in Approximate Addition for Low-Frequency Sound Zones. In *2022 18th Biennial Baltic Electronics Conference, BEC 2022* Article 9935613 IEEE (Institute of Electrical and Electronics Engineers). https://doi.org/10.1109/BEC56180.2022.9935613

General rights
Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal -

Take down policy
If you believe that this document breaches copyright please contact us at vbn@aub.aau.dk providing details, and we will remove access to the work immediately and investigate your claim.

Downloaded from vbn.aau.dk on: December 05, 2025

Error Reduced Carry Prediction in Approximate Addition for Low-Frequency Sound Zones

Peter Koch
Department of Electronic Systems
Aalborg University
Aalborg, Denmark
pk@es.aau.dk

Jan Østergaard
Department of Electronic Systems
Aalborg University
Aalborg, Denmark
jo@es.aau.dk

***Abstract*—Generation of personal listening spaces and sound zones is an emerging smart technology which requires multiple loudspeakers with individual pre-processing FIR control filters in order to establish separate sound experiences in enclosed spaces. Earlier studies have shown that for low frequencies, the acoustic contrast ratio between the individual sound zones is more sensitive to the word-length of the adder than to the word-length of the multiplier in a fixed-point implementation of the FIR filter operations.**

In this work, we investigate a possible hardware complexity reduction of the adder circuitry by replacing the normal accurate fixed-point adder with an approximate adder based on error reduced carry prediction. Compared to a total word-length of 12 bits for the adder, our study reveals that the hardware complexity can be reduced by 8% by replacing the two least significant bits of the accurate part of the addition with inexact computations. The resulting reduction in acoustic contrast ratio is less than 0.1 dB. For higher reference word-lengths such as 14 bits, the reduction in hardware complexity increases to 17% at a negligible reduction in acoustic contrast ratio of about 0.05 dB.

***Index Terms*—Sound Zones, Fixed-Point, Approximate Addition, Simulation, Hardware Complexity Reduction**

I. Introduction

The combination of multiple loudspeakers and advanced signal processing techniques makes it possible to control the sound field in a room and generate personal listening spaces or individual sound zones [1]. This is a smart technology which is currently being developed, e.g., for implementation of sophisticated soundscapes in private residences or in public locations such as hospitals. It is common to generate two sound zones; a bright zone and a dark zone [1], [2]. In the bright zone, the aim is to obtain a relatively high sound pressure level and a high sound quality, whereas in the dark zone the aim is to minimize the sound pressure and ideally obtain a silent zone. An efficient technique for generating sound zones is to design finite impulse response (FIR) pre-processing control filters for the individual loudspeakers (sub-woofers) taking into consideration the room impulse response in order to control the low-frequency part of the spectrum, typically up to 600 Hz, [1]–[7]. The mid-range and high-frequency part of the spectrum is normally controlled by traditional beam forming using a set of constant directivity loudspeakers and associated control filters which however, do not consider the room characteristics.

The computational complexity of sound zone generation therefore scales with the number and the length of the FIR control filters. For low-frequency sound zones, which is our focus in this work, the length of the filters is typically more than 100 taps, and one needs at least 8 loudspeakers to have sufficient low-frequency control of the sound field in order to generate two sound zones [8]. To reduce the complexity of the control filter computations, it is possible to use fixed-point implementation of the arithmetic operations [9]. In real-time scenarios, a fixed-point number representation is often applied in terms of 2's complement encoding of signals and coefficients. This is particularly true for dedicated architectures implemented in a configurable hardware technology, e.g., Field Programmable Gate Arrays, where the word-length can be adapted to various numerical needs. In order to prepare the sound zone application to be conducted in real-time, we demonstrate in [10] that the degradation of the acoustic contrast ratio (i.e., the ratio of the sound pressure levels in the bright and dark zones) can be kept less than 0.1 dB when the word-length is reduced to 12 bits for additions, and 8 bits for multiplications in the FIR filters.

Based on [10], we therefore conclude that the contrast ratio is significantly more word-length sensitive as related to additions than to multiplications. Seen from a hardware complexity point of view, this is a very interesting and convenient result due to the normally high complexity of multiplier circuits. At the same time however, since the adder word-length is the numerically limiting factor, it is most relevant to consider alternative methods for potential adder complexity reductions.

In this work we therefore investigate to what extend we can introduce a hardware complexity minimization of the adder circuitry subject to no or limited acoustic contrast ratio degradation. Basically, we evaluate the possibilities of employing approximate addition as an alternative to exact addition. We apply a recently published approximate adder design denoted Error Reduced Carry Prediction Approximate Adder (ERCPAA), [11]. It has better or similar performance than a variety of other approximate adders in terms of metrics such as i) carry prediction error rate, ii) mean relative error distance, and iii) normalized mean error distance as compared to power, energy and area-delay product. Our study reveals that for a total adder word-length of 12 bits, one can reduce

the hardware complexity by 8% by decreasing the accuracy by only 2 bits, thus introducing approximate computing in the filtering process. The resulting reduction in acoustic contrast ratio is less than 0.1 dB.

The rest of the paper is organized with a definition in section II of the mathematics behind spatial sound pressures, and the quality metrics used to quantify it. Section III explains the fundamentals of approximate addition, and next in section IV we elaborate on how we combine approximate addition into a sound zone simulation environment, including a description of the ERCPAA circuit. We describe our experiments in section V, elaborate on the results achieved in section VI, and finally draw the conclusions in section VII.

II. Sound Zones – Definitions and Metrics

We use L loudspeakers to generate two sound zones; a bright zone and a dark zone. In the bright zone, we measure the sound pressure level $p_b^{(m)}[k]$ at time k using microphone m, where $m = 1, \ldots, M_b$, and M_b denotes the number of microphones. The sound pressure $p_b^{(m)}[k]$ can be modelled by a linear convolution of the audio signal with the sound zone control filters and the room impulse responses (RIRs), respectively, i.e., [8],

$$p_b^{(m)}[k] = \sum_{\ell=1}^{L} \left(\bar{h}_b^{(m,\ell)} \star \bar{w}^{(\ell)} \star u \right)[k] \tag{1}$$

$$= \sum_{\ell=1}^{L} \sum_{j=0}^{N_w-1} \sum_{i=0}^{N_h-1} \bar{h}_b^{(m,\ell)}[i] \bar{w}^{(\ell)}[j] u[k-i-j] \tag{2}$$

where $\bar{h}_b^{(m,\ell)}$ denotes the RIR from the ℓ'th loudspeaker to the m'th microphone in the bright zone, $\bar{w}^{(\ell)}$ denotes the control filter for the l'th loudspeaker, and u is the audio signal. We define the sound pressure $p_d^{(m)}[k]$ in the dark zone in a similar manner, i.e.,

$$p_d^{(m)}[k] = \sum_{\ell=1}^{L} \sum_{j=0}^{N_w-1} \sum_{i=0}^{N_h-1} \bar{h}_d^{(m,\ell)}[i] \bar{w}^{(\ell)}[j] u[k-i-j]. \tag{3}$$

The average accumulated squared sound pressure levels P_{bright} and P_{dark} in the bright and dark zones, respectively, are given by,

$$P_{\text{bright}} \triangleq \frac{1}{N_u M_b} \sum_{m=1}^{M_b} \sum_{k=0}^{N_u-1} |p_b^{(m)}[k]|^2 \tag{4}$$

$$P_{\text{dark}} \triangleq \frac{1}{N_u M_d} \sum_{m=1}^{M_d} \sum_{k=0}^{N_u-1} |p_d^{(m)}[k]|^2 \tag{5}$$

where N_u denotes the length of the time-domain audio signal $\{u[k]\}$. The mean acoustic contrast ratio (expressed in dB) is defined as follows [2],

$$C = 10 \log_{10} \left(\frac{P_{\text{bright}}}{P_{\text{dark}}} \right) [\text{dB}]. \tag{6}$$

Let $\tilde{p}_b^{(m)}[k]$ be a specific desired target pressure level at time k at the m'th microphone position in the bright zone. Then the normalized mean-square error (nMSE) is defined as [2],

$$Q = 10 \log_{10} \left(\frac{\sum_{m=1}^{M_b} \sum_{k=0}^{N_u-1} |p_b^{(m)}[k] - \tilde{p}_b^{(m)}[k]|^2}{\sum_{m=1}^{M_b} \sum_{k=0}^{N_u-1} |\tilde{p}_b^{(m)}[k]|^2} \right) [\text{dB}]. \tag{7}$$

III. Approximate Addition

A traditional 2's complement d-bit Ripple Carry Adder (RCA) can be opted for when exact additions are requested. Using such an adder however, introduces a carry chain with complexity $\mathcal{O}(d)$. There exists a vast amount of alternatives to the RCA which, still conducting the addition accurately, can break the carry chain and thus reduce the time complexity, e.g., Carry Look-Ahead and Carry Select adders. Such alternative adders come with an increased hardware complexity due to advanced carry prediction circuits.

However, *approximate computing* introduces the possibility for simultaneous reduction of the execution time and the area overhead at the expense of a reduced numerical accuracy, [12]. For instance, an approximate adder can be constructed from an RCA which is subdivided into two sections; i) an MSB portion with normal carry propagation which performs a standard accurate addition, and ii) an approximate LSB section where the bit-wise adders (i.e., Full Adders (FA)) are replaced with more simple adder elements (e.g., OR-gates). At the same time, the carry propagation, if any at all, is significantly simplified. In most of such adders, there is no carry propagation internally in the approximate part, except from a single speculative carry which is generated in the MSB end and fed into the accurate part in order to initiate a normal carry propagation process. Due to these two characteristics, the inaccurate LSB section is typically less power consuming, and similarly characterized by a reduced overall propagation delay as compared to an accurate adder. The benefit, as compared to e.g., a standard RCA, is a reduced power-delay product.

IV. Approximate Addition Used in Sound Zones

Sound zone generation includes two parts; 1) the audio signal $\{u[k]\}$ is convoluted with the impulse responses of the control filters $\{\bar{w}^{(\ell)}\}$ (i.e., the coefficients in the FIR filters) to obtain the control signals which are then 2) played out by the loudspeakers and thereby convoluted with the impulse response of the room, see Equ. (1) and (3). Focusing only on the first part, we modify the filter coefficients $\{\bar{w}^{(\ell)}\}$ as well as the convolution operator $\star$ in order to model the effect of accurate finite-precision multiplications and inaccurate finite-precision additions in the control-filter multiply-accumulate operations.

We design the control filters using the method outlined in [3]. The filter coefficients are then scaled so that they are within the range $]-1, 1]$. Similarly, the audio signal is scaled to the range $]-1/3, 1/3]$, thus eliminating the effect of overflow in the multiply-accumulate operations. For the multiplications we build a simulation model of a 2's complement accurate

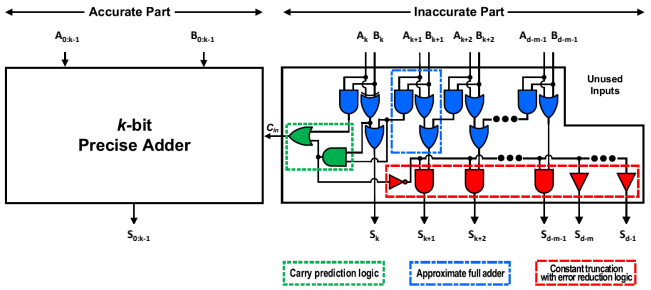

Fig. 1. A d bit ERPCAA consist of an accurate k bit MSB part, and an inaccurate $d-k$ bit LSB part which is further subdivided into an m bit unused section at the LSB end (here all bits are either "0" or "1"), and a $d-k-m$ bit approximate section located towards the MSB end. A predicted carry calculated in the inaccurate part is fed into the accurate part. Courtesy of [11] (note however that we use bit indexing from 0 to $d-1$ which is opposite to the original work where the bits are indexed $n-1$ down to 0).

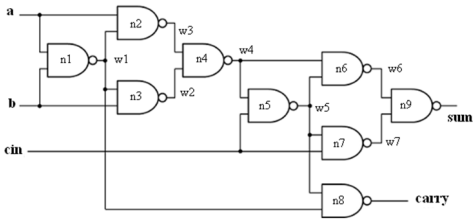

Fig. 2. The 2-input XOR gate applied for generating both the FA sum and the FA carry has a CMOS transistor equivalent which is significantly more complex than corresponding 2-input AND, OR, and NAND gates. Therefore, in order to make a fair comparison of the gate saving achieved when replacing true FAs with approximate FAs, we compare against this 9-gate 2-input NAND-gate implementation of an FA.

fixed-point multiplier with a word-length of 8 bits as suggested in [10]. For the additions, we build a parameterized simulation model of the ERCPAA approximate adder. The circuit diagram of this adder is shown in Fig. 1

The inaccurate part of the adder consists of three basic functions denoted i) approximate full adder (blue), ii) carry prediction logic (green), and iii) constant truncation with error reduction logic (red). The approximate FA element replaces the accurate FA which have three inputs a, b, and c_{in} and which functionally is characterized by the two Boolean equations

$$sum = (a \oplus b) \oplus c_{in} \tag{8}$$

$$carry = (a \cdot b) + c_{in} \cdot (a \oplus b). \tag{9}$$

Comparing these equations against Fig. 1, we first of all realize that the Boolean $a \oplus b$ operation in the sum calculation has been simplified to a Boolean $a + b$ function, i.e., an arithmetic addition is replaced by a logic addition, the exception is the most significant bit which remain being implemented as an exclusive OR of the two input operand bits. This is done in order not to degrade the accuracy at the MSB (the more significant a bit is, the more important is its numerical correctness). Next, the two terms of the $carry$ calculation, i.e., the carry generation and the carry propagation, have been reduced to a single carry generation, i.e., a Boolean $a \cdot b$ operation, thus eliminating the traditional longitudinal carry chain found in an RCA. Contrary to various other approximate adders which totally neglect carry propagation, the ERPCAA supports a simple, yet effective carry scheme which spans two consecutive bits, i.e., the carry bit generated at position j is logically added to the sum bit generated at position $j-1$.

In order to further improve the accuracy of the sum of the overall approximate adder, a carry is generated in the MSB end of the inaccurate part, and next fed into the accurate part as an ordinary $carry_{in}$ signal. This carry is based on i) the carries from the two most significant bits, and ii) the MSB sum. Using these three bits (contrary to using only a carry generate from the MSB), reduces the error in the carry prediction at the input to the accurate part, still supporting a significantly lower gate count as compared to a full size RCA.

Finally, the ERPCAA introduces a mechanism to control all the sum bits from bit position $k+1$ down to $d-1$, the control signal being calculated as $\overline{sum_k \cdot carry_{k+1}}$ (from the approximate FAs), and therefore is equal to "1" when one or both of these bits equals "0". In the interval from $k+1$ down to $d-m-1$, the control signal, when equal to "1", is used to pass the OR-combined sum- and carry-bits from the approximate FAs directly to the final sum output bits, s_i. On the other hand, a control signal value equal to "0" forces all sum bit in this specific interval equal to "0" in order to numerically compensate for the s_k bit which in this particular situation equals a faulty "1". The control signal is fed directly to the sum output bits, s_i, from bit position $d-m$ down to $d-1$ thereby generating an m-bit constant output and thus eliminating the need for approximate FAs in the LSB end of the inaccurate part.

In our parameterized implementation, the overall word-length d, the length of the accurate part k, and the length of the constant section m can all be tuned individually.

V. Experimental Setup

In order to enable a comparison of the ERPCAA against a traditional RCA, considering not only the computational accuracy, we investigate the hardware complexity of the two different adder approaches. Normally, one would (in an ASIC design) implement an RCA using FAs based on Exclusive-OR gates as outlined in Equ. (8) and (9). In our context however, where we want to evaluate the potential hardware complexity reduction in terms of *gate count savings* when substituting RCA-based addition with an ERPCAA, it would lead to an unfair comparison unless we consider the RCA being implemented with 2-input gates having complexities comparable to the gates used in the inaccurate part of the ERPCAA. The 2-input XOR gate has a CMOS transistor equivalent which is at least 50% more complex (in number of transistors) as compared to corresponding 2-input AND, OR, and NAND gates, [13], used in the inaccurate part of the approximate adder (these all consist of 2- and 1-input gates). For that reason, we realistically use the 9 gate 2-input NAND-

based FA realisation shown in Fig. 2 as the reference in our comparison.

We have performed a simulation study using real measured room impulse responses from a room of size 7.00x8.12x3.00 meters equipped with $L = 8$ loudspeakers. The "audio" signal is white noise sampled at 1200 Hz and with a frequency support from 20 Hz to 600 Hz. We average the performance over 12000 samples of the audio signal, and let $M_b = M_d = 20$ so that there are 20 microphones in each of the two sound zones.

When using 32 bits fixed-point adders and multipliers (corresponding to "infinite" word-length) the contrast ratio (6) is 23.09 dB and the nMSE (7) is -9.89 dB. Onward, these values are considered as references. Using our ERPCAA simulation model, i) the resulting contrast ratio, ii) the nMSE, and iii) the gate count saving (absolute number/percentage as compared to a d-bit RCA) when experimenting with different d-, k-, and m-values are shown in Tables I and II. The tables represent different total word-lengths for the ERPCAA, $d \in [6; 20]$, evaluated in steps of 2 bits. For all simulations, we apply a multiplier with a word-length equal to 8 bit, 2's complement fixed-point.

For a given total word-length d, and a given word-length of the accurate part k, we experiment with the word-length of the constant section m, i.e., $(d-k-m)$ represents the section of the inaccurate part where FAs are replaced by approximate FAs. This section of the adder, if it exists, should always have a length of at least 2 bits, i.e., bit index k and $k+1$, because the carry to the accurate part, i.e., c_{in} in Fig. 1, is generated based solely on these two bits. For $k = d$, we force this carry equal to "0", and similarly for this situation we notice that there are no savings in the hardware complexity. Finally, we note that those entries in Tables I and II which are left blank represent settings which are not legal.

TABLE I
CONTRAST RATIO (dB), nMSE (dB), AND GATE COUNT SAVING (ABSOLUTE NUMBER/PERCENTAGE) FOR $d = 6, \cdots, 14$.

d=6	$(d-k-m), m$									
k	0,0	2,0	4,0	3,1	2,2	6,0	5,1	4,2	3,3	2,4
6	14.95 -8.25 0/0									

d=8	$(d-k-m), m$									
k	0,0	2,0	4,0	3,1	2,2	6,0	5,1	4,2	3,3	2,4
8	20.67 -9.69 0/0									
6		14.44 -8.73 9/13								

d=10	$(d-k-m), m$									
k	0,0	2,0	4,0	3,1	2,2	6,0	5,1	4,2	3,3	2,4
10	22.72 -9.86 0/0									
8		21.35 -9.75 9/10								
6			18.00 -9.33 19/21	17.67 -9.26 22/24	15.68 -8.90 25/27					

d=12	$(d-k-m), m$									
k	0,0	2,0	4,0	3,1	2,2	6,0	5,1	4,2	3,3	2,4
12	22.98 -9.87 0/0									
10		22.88 -9.86 9/8								
8			22.11 -9.84 19/18	22.03 -9.83 22/20	21.42 -9.79 25/23					
6						18.02 -9.25 29/27	18.02 -9.25 32/30	17.97 -9.25 35/32	17.71 -9.20 38/35	15.73 -8.83 41/38

d=14	$(d-k-m), m$									
k	0,0	2,0	4,0	3,1	2,2	6,0	5,1	4,2	3,3	2,4
14	22.98 -9.87 0/0									
12		22.98 -9.87 9/7								
10			22.93 -9.87 19/15	22.92 -9.87 22/17	21.86 -9.86 25/20					
8						22.10 -9.85 29/23	22.10 -9.85 32/25	22.09 -9.85 35/28	22.00 -9.85 38/30	21.40 -9.80 41/33

VI. RESULTS AND DISCUSSION

Since essentially there are two conceptually different options for interpreting our results, we initially discuss this issue. We have defined the 12 bit addition (and the 8 bit multiplication) as an acceptable fixed-point word-length because the contrast ratio can be kept within 0.1 dB as compared to the infinite word-length reference. Introducing approximate computing in term of the ERPCAA, one may either consider this adder an opportunity to i) add more approximate bits while maintaining the fundamental accurate word-length (k), thus improving the performance with only limited increased hardware overhead, or ii) experiment with different total word-lengths (d) and then observe how performance and hardware overhead is impacted by changing values of k and m.

Our prime objective with this study is to evaluate how approximate addition can be introduced into the low-frequency sound zone application, and therefore we have no preferences towards which one of the two mentioned approaches to use – both are equally applicable in their own right.

We therefore consider the results from a general perspective and report on the following findings which are all extracted from the numbers provided in Tables I and II;

- For maintained total word-length d, we see a performance increase (contrast ratio and nMSE) when the word-length of the accurate part is increased, no matter the length of the inaccurate part. Under the same conditions, we observe an expected decrease in the hardware saving.
- For maintained total word-length d, and maintained

TABLE II
CONTRAST RATIO (dB), nMSE (dB), AND GATE COUNT SAVING (ABSOLUTE NUMBER/PERCENTAGE) FOR $d = 16, \cdots, 20$.

d=16	$(d-k-m), m$									
k	0,0	2,0	4,0	3,1	2,2	6,0	5,1	4,2	3,3	2,4
14		23.04 -9.87 9/6								
12			22.97 -9.87 19/13	22.98 -9.87 22/15	22.97 -9.87 25/17					
10						22.92 -9.86 29/20	22.92 -9.86 32/22	22.92 -9.86 35/24	22.91 -9.86 38/26	22.85 -9.86 41/28

d=18	$(d-k-m), m$									
k	0,0	2,0	4,0	3,1	2,2	6,0	5,1	4,2	3,3	2,4
14			23.04 -9.87 19/12	23.04 -9.87 22/14	23.04 -9.87 25/15					
12						22.98 -9.87 29/18	22.98 -9.87 32/20	22.98 -9.87 35/22	22.98 -9.87 38/23	22.97 -9.87 41/25

d=20	$(d-k-m), m$									
k	0,0	2,0	4,0	3,1	2,2	6,0	5,1	4,2	3,3	2,4
14						23.04 -9.87 29/16	23.04 -9.87 32/18	23.04 -9.87 35/19	23.04 -9.87 38/21	23.04 -9.87 41/23

length of the accurate part k, we see a performance increase when the length of the constant section m is decreased. Under the same conditions, we always observe a decrease in the hardware saving.

- The percentage performance decrease observed for an increase in the constant section m is reduced for increasing total word-length d.
- For a fixed word-length k of the accurate part, the hardware saving percentage increases with increased total word-length d.

More specifically, we observe how the reference performance, i.e., a 22.98 dB contrast ratio and a -9.87 dB nMSE can be realized for different settings of the variable word-length parameters;

- It takes at least a total word-length $d = 12$ in order to establish the reference performance.
- An 8% hardware saving can be achieved by reducing the accurate part with 2 bit. In such a case the nMSE essentially remains unchanged while the contrast is further decreased with 0.1 dB.
- The reference can be achieved exactly for $d = 14$ and either $k = 14$ or $k = 12$. In the latter case, the hardware saving is 7% as compared to the first case. In both cases, the nMSE stays the same. Further decreasing k to 10 and for $m = 1$, the contrast ration is reduced by only 0.06 dB and the nMSE is unchanged with a hardware reduction equal to 17% as compared to an accurate 14 bit RCA.
- Increasing d to 16 and beyond, it is always possible to establish the reference performance (in some cases it is even increased due to the extended word-length).
- Decreasing d to 10, essentially it is possible to maintain the nMSE for an accurate adder while the contrast is decreased with 0.26 dB. Introducing a 2-bit inaccurate part, the contrast is deteriorated with 1.63 dB and the nMSE is down to -9.75 dB at the benefit of a 10% hardware saving as compared to an accurate 10 bit RCA.

VII. CONCLUSIONS

We have investigated how the complexity of a real-time hardware architecture can be reduced using approximate addition, still maintaining the same or only insignificant reduced performance of the sound field in a low-frequency sound zone environment. For the recently proposed ERPCAA approximate adder, which in terms of numerical capabilities is superior or comparable to most other approximate adders, we demonstrate that for a 12-bit reference word-length, an 8% reduction in the total gate count for the adder circuitry is immediate possible in return for an acceptable 0.1 dB degradation of the acoustic contrast between the bright and dark zones. From a general perspective, we have clearly demonstrated that approximate addition is a powerful approach towards hardware complexity reduction. Particularly, we emphasize the possibility not only to adjust the total adder word-length, but also to consider the word-length of i) the accurate part of the adder, and ii) the constant section to be valuable and interesting parameters which can be tuned in order to devise a solution which balances audio performance and hardware usage.

REFERENCES

[1] S. J. Elliott, J. Cheer, H. Murfet, and K. R. Holland, "Minimally radiating sources for personal audio," *The Journal of the Acoustical Society of America*, vol. 128, 2010.

[2] M. Gálvez, S. J. Elliott, and J. Cheer, "Time domain optimization of filters used in a loudspeaker array for personal audio," *IEEE/ACM Transactions on Audio Speech and Language Processing*, vol. 23, pp. 1869 – 1878, 2015.

[3] M. Møller and M. Olsen, "Sound zones: On envelope shaping of FIR filters," *24th Int. Congress on Sound and Vibration, ICSV 2017*, 2017.

[4] Y. Cai, M. Wu, L. Liu, and J. Yang, "Time-domain acoustic contrast control design with response differential constraint in personal audio systems," *The Journal of the Acoustical Soc. of America*, vol. 135, 2014.

[5] Y. Cai, M. Wu, and J. Yang, "Design of a time-domain acoustic contrast control for broadband input signals in personal audio systems," 2013.

[6] A. Canclini, D. Markovic, M. Schneider, F. Antonacci, E. A. Habets, A. Walther, and A. Sarti, "A weighted least squares beam shaping technique for sound field control," vol. 2018-April, 2018.

[7] M. B. Møller and J. Østergaard, "A moving horizon framework for sound zones," *IEEE/ACM Transactions on Audio Speech and Language Processing*, vol. 28, 2020.

[8] M. Møller, "Sound zone control inside spatially confined regions in acoustic enclosures," *PhD Thesis, Aalborg University*, 2019.

[9] M. Aktan, A. Yurdakul, and G. Dundar, "An algorithm for the design of low-power hardware-efficient fir filters," *IEEE Trans. Circuits Syst.*, vol. 55, pp. 1536 – 1545, 2014.

[10] P. Koch and J. Østergaard, "The effect of fixed-point arithmetic on low frequency sound-zone control," *Proceedings of The Baltic Nordic-Acoustic Meetings '22*, Aalborg, Denmark, May 2022.

[11] J. Lee, H. Seo, H. Seok, and Y. Kim, "A novel approximate adder design using error reduced carry prediction and constant truncation," *IEEE Access*, vol. 9, pp. 939–953, Sep. 2021.

[12] H. Jiang, F. J. H. Santiago, H. Mo, L. Liu, and J. Han, "Approximate arithmetic circuits: A survey, characterization, and recent applications," *Proceedings of The IEEE*, vol. 108, pp. 2108 – 2135, 2020.

[13] N. Weste and K. Eshraghian, *Principles of CMOS VLSI Design - A Systems Perspective*. Addison-Wesley Publishing Company, 1985.