Low-precision Logarithmic Number Systems: Beyond Base-2

SYED ASAD ALAM, JAMES GARLAND, and DAVID GREGG, Lero,
Trinity College Dublin, Ireland

Logarithmic number systems (LNS) are used to represent real numbers in many applications using a constant base raised to a fixed-point exponent making its distribution exponential. This greatly simplifies hardware multiply, divide, and square root. LNS with base-2 is most common, but in this article, we show that for low-precision LNS the choice of base has a significant impact.

We make four main contributions. First, LNS is not closed under addition and subtraction, so the result is approximate. We show that choosing a suitable base can manipulate the distribution to reduce the average error. Second, we show that low-precision LNS addition and subtraction can be implemented efficiently in logic rather than commonly used ROM lookup tables, the complexity of which can be reduced by an appropriate choice of base. A similar effect is shown where the result of arithmetic has greater precision than the input. Third, where input data from external sources is not expected to be in LNS, we can reduce the conversion error by selecting a LNS base to match the expected distribution of the input. Thus, there is no one base that gives the global optimum, and base selection is a trade-off between different factors. Fourth, we show that circuits realized in LNS require lower area and power consumption for short word lengths.

CCS Concepts: • **Computer systems organization** → **Embedded systems**; • **Hardware** → Integrated circuits; • **Mathematics of computing**;

Additional Key Words and Phrases: Logarithmic number system; quantization error; digital arithmetic

ACM Reference format:

Syed Asad Alam, James Garland, and David Gregg. 2021. Low-precision Logarithmic Number Systems: Beyond Base-2. *ACM Trans. Arch. Code Optim.* 18, 4, Article 47 (July 2021), 25 pages.
<https://doi.org/10.1145/3461699>

This is a very significant extension and complete rewrite of a four-page work-in-progress (WIP) conference paper [1]. We add a discussion of LNS number distribution; the concept of base aliasing; more correct methods of rounding and computing rounding errors; the proposal and experimental evaluation of selecting the base to reduce the area and delay of LNS add/subtract circuits; a study of the impact the base on add/subtract units with larger output precision than input; experiments for word size in the range 5–16 as compared to just 8 in the WIP paper; implemented a number of circuits using logarithmic, fixed and floating number systems to show the advantage of LNS and many more details.

This material is based upon work supported, in part, by Science Foundation Ireland under Grant No. 13/RC/2094_P2 and, in part, by the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie Grant Agreement No. 754489. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author and do not necessarily reflect the views of Science Foundation Ireland and EU's Horizon 2020 programme.

Author's address: S. A. Alam, J. Garland, and D. Gregg, Lero, Trinity College Dublin, College Green, Dublin 2, Dublin, Ireland; emails: {syed.asad.alam, jgarland, dgregg}@tcd.ie.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2021 Association for Computing Machinery.

1544-3566/2021/07-ART47 \$15.00

<https://doi.org/10.1145/3461699>

1 INTRODUCTION

As the world moves toward autonomous systems in every field, there is a need for small embedded devices that can interact with the environment on its edge rather than offloading computations to the cloud. These devices often operate on physical-world data such as audio, video and input from various sensors, using computationally intensive techniques, like filtering, facial recognition, image segmentation, and object detection. Such techniques are based on signal processing or **deep neural networks (DNN)**. The ability to execute these applications on small resource-constrained embedded devices is a major milestones in truly enabling technologies like driverless cars, smart cities and smart mobile devices [34]. Given the increasing quantities of data processed by smart devices and autonomous systems, computing at the edge is essential. Sending data back to the cloud for processing requires large amounts of energy, and may become a bottleneck that fails to meet latency requirements for real-time embedded and autonomous systems while also stretching the limited processing, memory and/or battery capacity of embedded devices [12, 36].

One possible solution is a **logarithmic number system (LNS)**, which represents real numbers as a fixed-point exponent. Compared to floating point, LNS dramatically simplifies the hardware needed for multiplication, division, and square root. Multiplication and division become fixed-point addition and subtraction of the exponents, which can be implemented with a simple integer adder [2, 22]. Square root is computed by dividing the exponent by two, which can be implemented by simply dropping the least-significant bit of the fixed-point exponent. Among these basic operations, multiplication occur very frequently in systems like neural networks and **finite-impulse response (FIR)** filters and reduction of its complexity will have a major impact on the overall reduction of area and power. The downside of LNS is that addition and subtraction are non-trivial. A common strategy to implement LNS adders/subtractors use ROM look-up tables, which in the simple case are exponential in the size of LNS word. There is an extensive literature on compacting these tables for large word sizes [20, 26, 41].

Smaller word sizes can reduce the memory footprint of data, and the complexity of arithmetic for a variety of number systems [9, 14, 17, 33], which will have a significant impact on the ability to deploy systems in resource constrained embedded devices deployed on the edge of networks. Further, fixed-point number systems with very short word lengths have been proposed in the literature for a variety of signal processing applications [2, 24], while shorter floating- and fixed-point numbers have also been used for neural networks [8, 18, 19, 39, 43, 44].

An important aspect of any number system is the distribution of the representable numbers. LNS have a distribution that is non-uniform. Each number is a constant multiple of the previous one, resulting in an exponential distribution. Fully one half of LNS numbers fall between -1 and 1 , and representable numbers become increasingly sparse thereafter. In applications like FIR filters, the coefficient space lies within the range of ± 1 [2], while the distribution of weights and biases in DNNs is also typically clustered within a small range [25]. Thus, LNS is suitable for use in a number of different applications [2, 22, 30].

The exponential distribution of LNS depends on two factors: the base and the number of integer/fractional bits in the fixed-point exponent. In existing literature, the base is almost always a power of two, typically 2 , $\sqrt{2}$, or $\sqrt[3]{2}$. However, there is no systematic study of different bases in the literature, perhaps because within sensible ranges the choice of base has little impact on LNS with large word size. However, we show that for low-precision LNS the base has a major impact. We make the following contributions.

- We demonstrate that LNS with different bases have different arithmetic rounding errors, which is significant for low-precision LNS and can be reduced by selecting an appropriate base.

- We optimize hardware for low-precision LNS arithmetic by selecting a base with favourable add and subtract truth tables. We achieve large savings in circuit area and depth for LNS adder and subtractor circuits by implementing the tables using logic gates as compared to ROM-based implementation.
- We show that where input data must be converted from another number format to LNS, base selection can enable a more accurate approximation of the original data distribution, where accuracy is measured by the representation error.
- We also define arithmetic and representation error in a multiplicative sense in the real domain, which allows us to describe the error in bits.
- We analyze mixed-precision tables where the values stored in addition and subtraction tables have more precision than their inputs and show that realizing these tables results in a reduced error without an exponential increase in the hardware cost to realize them.
- We synthesis adders, subtractors and FIR filters for different word lengths using different number systems and show that circuits LNS-based circuits achieve lower area and power consumption as compared to its counterparts for short word lengths.

The article is organized as follows: Section 2 describes related work. Section 3 gives a brief introduction of logarithmic number system, the effect of base aliasing and of rounding of numbers in LNS. Section 4 discusses the quantization of floating-point numbers into LNS. Section 5 describes LNS addition and subtraction and approximation errors in these two operations. The synthesis of addition and subtraction tables is described and evaluated in Section 6, while Section 7 presents a discussion on the trade-offs in selecting the best base based on various parameters. Mixed-precision tables are the focus in Section 8, while Section 9 discusses the hardware cost of realizing LNS adder/subtractor circuits and their comparison against fixed- and floating-point circuits. The discussion of hardware cost of real word circuits is extended in Section 10 for FIR filters using the different number systems and word length.

2 RELATED WORK

LNS has been evaluated as an alternative to fixed- and floating-point numbers in a number of applications. It provides better dynamic range and round-off noise performance as compared to fixed point [21] and floating point [4] in some signal processing applications. Paliouras and Stouratis [32] show that LNS can lead to lower switching activity leading to reduced power consumption while also reducing representational error relative to fixed point. Basetas et al. [6] show that LNS requires a reduced word length when implementing FIR and infinite impulse response filters. Alam and Gustafsson have also shown [2] that LNS improves the approximation error by around 20% when designing FIR filters as compared to a fixed-point design. Coleman proposed a 32-bit logarithmic processor and concludes that LNS provides faster execution, more accuracy and reduced architectural complexity as compared to single-precision 32-bit floating point, provided that the adder/subtractor has a latency less than that of a 32×32 -bit multiplier. Chandra [10] shows an order of magnitude improvement in error-to-signal variances for round-off noise in FIR filters using LNS over floating-point number representation.

LNS has also been explored for deep neural networks using very short word lengths. The LogNet inference engine of Lee et al. [25], uses 4- and 5-bit LNS to represent weights and improves results as compared to fixed-point encoding without retraining while keeping the activations in floating point. They also encode the activations in LNS with 3- and 4-bits while keeping the weights in floating point and show negligible to zero degradation in performance as compared to floating point. For example, the top-5 accuracies (without training) using logarithmic encoding in base-2 with 4 and 5 bits is 73.4% and 74.6%, respectively, for AlexNet and 85.2% and 86.0%, respectively, for

VGG16. The corresponding accuracy using 32-bit floating-point numbers is 78.3% and 89.8% for the two nets, respectively, and much better than corresponding fixed-point performance. This number further improves to a maximum of 85.2% upon training and logarithmic encoding also has more sparse weights (35.8% against 21.2%). However, although not clearly stated, the exponent for LNS is represented using integers only and the impact of using a fixed-point exponent is not explored.

Miyashita et al. also use an integer representation for the exponent while quantizing using 3- and 4-bit encodings for activations and 5-bit for weights [30]. For AlexNet, Miyashita et al. show that a 4-bit encoding achieves only a 1.4% drop in the top-5 accuracy for AlexNet and no drop for VGG16. The corresponding drop in accuracy for 3-bit logarithmic encoding is also negligible. Miyashita et al. also report a reduction in the memory requirements for AlexNet and VGG16 using LNS encoding of 85.7% and 77.9%, respectively. They have also considered bases of $\sqrt{2}$ and $\sqrt[3]{2}$ for representation and find that the total quantization error on the weights for $\sqrt{2}$ is about $2\times$ smaller than base-2 LNS. However, as we show in Section 3.1, bases 2, $\sqrt{2}$ and $\sqrt[3]{2}$ are simple aliases of one another, where the binary point of the fixed-point exponent is moved left. Vogel et al. [42], present a low word length-based quantization method for LNS to be used in a neural network and have shown LNS to achieve 22.3% lower power as compared to an 8-bit fixed-point-based design. Arnold et al. [3] present an approach to implement back propagation using tableless LNS ALU based on modified Mitchell's method [28] and achieve one third reduction in the hardware resources as compared to a conventional fixed-point implementation.

Sicuranza [37] explored varying the LNS base on the accuracy of a 2D digital filter. Dimitrov et al. [16] proposed an LNS that uses two different integer LNS bases. However, to the best of our knowledge, a comprehensive and systematic study of LNSs with respect to different bases and error measurement is missing in the literature. Also missing is the study of the impact of choice of base on short word length systems and providing a study of the trade-off involved in selecting bases as a function of the word length while also analyzing the impact of LNS on systems like FIR filters, which are heavily dependent on the operations of multiplications and additions. These important aspects form the bulk of the contribution of this work.

3 LOGARITHMIC NUMBER SYSTEM (LNS)

A number in the logarithmic number system is represented using the zero-bit, sign-bit and the actual logarithmic value, represented as a triplet $(\mathcal{X})$, as follows [40]:

$$\mathcal{X} = (z_x, s_x, m_x), \quad (1)$$

where z_x indicates whether $\mathcal{X}$ is zero, $s_x = \text{sign}(X)$ and $m_x = \log_b |X|$. The exponent (or logarithmic) value m_x is typically represented using two's complement fixed-point representation, with i and f integer and fractional bits (mathematically represented as $Q(i, f)$) [2, 41]. The base- b of an LNS number influences the representation capabilities and complexity of computations and conversions [38]. The real value zero cannot be expressed in the form b^m , $b \neq 0$, and therefore LNS commonly uses an additional bit to represent zero. Some LNS instead use the most negative possible m_x as a special value to represent zero [20].

3.1 Bases and Base Aliasing in LNS

The base b is a key parameter of the LNS. For a base $b > 1$, each LNS value is exponentially larger than the previous one. However, the ratio between successive LNS values is determined not just by the base but also by the number of fractional bits in the fixed-point format $Q(i, f)$. The least significant bit of the fixed-point exponent has the value 2^{-f} in the *log domain* of the LNS number, which we define as the *unit of least-precision* in the log domain (ULP_{LNS}). Increasing the log domain value by one ULP_{LNS} corresponds to multiplying the LNS number by $b^{2^{-f}}$ in the *real domain*, which

we simply define as ULP. We further define an integer radix form, r of the LNS base, where $r = b^{2^{-f}}$. The ratio between successive values of the LNS in the real domain is r .

Note that the LNS with $Q(i, f)$ and base b contains the same set of values as the LNS with $Q(i + f, 0)$ and base r . Thus, the effect of changing the base of an LNS may be equivalent to simply moving the binary point in the fixed-point representation of m_x of Equation (1). This is shown for base-2 and $\sqrt{2}$ in Equation (2). Through simple manipulation of powers, Equation (2) shows that a number represented using base- $\sqrt{2}$ and three integer and four fractional bits for the exponent is equivalent to the number represented using base-2 and two integer and five fractional bits:

$$\sqrt{2}^{i_2 i_1 i_0 . f_1 f_2 f_3 f_4} = 2^{\frac{i_2 i_1 i_0 . f_1 f_2 f_3 f_4}{2}} = 2^{i_1 i_0 . f_1 f_2 f_3 f_4}. \quad (2)$$

In other words, LNS with base- $\sqrt{2}$ and exponent of the form $Q(i, f)$ contains exactly the same set of values as base-2 LNS with $Q(i - 1, f + 1)$, and base- $\sqrt[4]{2}$ with $Q(i + 1, f - 1)$. By selecting the position of the fixed point, standard base-2 LNS can represent implicit bases such as $\sqrt[3]{2}$, $\sqrt[4]{2}$, $\sqrt[5]{2}$, 2, 4, 8, and so on. Remarkably, despite a very extensive literature on LNS over several decades, we have been unable to find a statement of this simple identity in the literature. For example, Miyashita et al. [30] explored the effect of varying the LNS base by experimentally evaluating the bases 2, $\sqrt{2}$ and $\sqrt[4]{2}$. By exploring only bases that can be already represented using base-2, Miyashita et al. [30] investigated only the effect of moving the binary point in m_x , not in exploring arbitrary bases.

Thus, when choosing a base for the LNS, the important factors are the total number of bits in m_x , that is $i + f$, and the integer radix form of the base, r . In principle, fixed-point exponents are entirely unnecessary, and we can simply use an integer exponent and a suitable base $b = r$. However, practical values of r are typically not very convenient numbers to refer to. Therefore, we maintain the existing tradition of representing m_x as a fixed-point number, rather than an integer. We explore bases in the range $\sqrt{2}$ to 2, which covers a range of values separated by one binary point position in the format of m_x .

3.2 Rounding LNS Numbers

LNS values have a finite number of digits, and therefore cannot represent all real numbers exactly. Where a real number arising from arithmetic cannot be represented exactly in the LNS, we normally round to a nearest representable LNS number. Rounding to nearest is a common strategy in floating-point number systems; in non-exceptional cases the FP rounding error is bounded by half of one FP mantissa unit of least precision.

A problem with rounding to nearest in LNS is that there are two different and incompatible possible approaches. One option is to round in the LNS domain, i.e., to round the LNS exponent to the nearest representable fixed-point exponent m_x . Another option is to round in the real domain, that is to consider the LNS number and its nearest neighbours to the real domain, and choose the rounding that minimizes the real domain error.

The two approaches give slightly different answers. Consider the case of rounding the fixed-point *binary* exponent value 0.1001 in a base-2 LNS with only integer digits. The choice is between rounding this value downwards to an exponent of 0, or upwards to an exponent of 1. If we round in the log domain, then it is clear that 0.1001 is greater than the mid-point of 0.1, and we should therefore round upwards to 1. However, if we convert these numbers to the real domain, then we are choosing between rounding downwards to $2^0 = 1$ or upwards to $2^1 = 2$. The value of binary fixed-point exponent 0.1001 is $2^{\frac{9}{16}} \approx 1.44043$, which is clearly closer to 1 than 2. Thus, if we round in the real domain, then we should round this number downwards.

In addition to rounding in the real or log domain, one needs to decide whether to measure absolute or relative error. For fixed-point numbers, one either uses the absolute error between the

original value and the rounded one (Equation (3)) or the relative error (Equation (4)), as given here:

$$e_a = \hat{m}_x - m_x, \quad (3)$$

$$e_r = \frac{e_a}{m_x}, \quad (4)$$

where e_a and e_r represent the absolute and relative error, respectively.

The measure of error used in literature has typically been the relative error but in the real domain [11], as given by

$$e_{rr} = \frac{b^{\hat{m}_x} - b^{m_x}}{b^{m_x}}, \quad (5)$$

where e_{rr} is the relative error in the real domain and b is the base. However, recall that in LNS adding one ULP_{LNS} to a number in the log domain corresponds to multiplying the real domain value by the integer radix form, r , of the base, given by $b^{2^{-f}}$ and defined as ULP in Section 3.1. Thus, rounding in the log domain minimizes the *multiplicative error* of the real domain value, that is the factor multiple between the exact real-domain value and the rounded real domain value. However, rounding in the real domain minimizes the real domain *additive relative error*, as given in Equation (5), that is the amount that must be added to or subtracted from the exact real domain value to get the rounded real domain value. We define the multiplicative error in the real domain as

$$e_m = \frac{b^{\hat{m}_x}}{b^{m_x}} = b^{\hat{m}_x - m_x}, \quad (6)$$

where b is the base, as usual. The range of this error, ignoring round to zeros, is given by

$$\frac{1}{\text{ULP}_h} \leq e_m \leq \text{ULP}_h, \quad (7)$$

where $\text{ULP} = b^{2^{-f}}$ and $\text{ULP}_h = b^{2^{-f}-1}$ with b being the base and f being the number of fractional bits. Since this error measure is based on a ratio, it makes more sense to calculate the geometric mean of it, as given by

$$\mu_g = \left(\prod_{i=1}^n e_{m,i} \right)^{\frac{1}{n}}, \quad (8)$$

where $e_{m,i}$ is the individual multiplicative relative error in the real domain given by Equation (6) and n is the total number of values over which the mean is calculated.

The geometric mean corresponds to the exponential of the arithmetic mean of logarithms given by

$$\begin{aligned} \left(\prod_{i=1}^n e_{m,i} \right)^{\frac{1}{n}} &= b^{\left[\frac{1}{n} \sum_{i=1}^n \log_b(e_{m,i}) \right]} \\ &= b^{\left[\frac{1}{n} \sum_{i=1}^n (\log_b(b^{\hat{m}_{x,i} - m_{x,i}})) \right]} \\ &= b^{\left[\frac{1}{n} \sum_{i=1}^n \hat{m}_{x,i} - m_{x,i} \right]}, \end{aligned} \quad (9)$$

where $e_{m,i}$ is given by Equation (8), b is the base, $\hat{m}_x$ and m_x are the rounded and original exponent (as given in Equation (1)). Here, we take the logarithm of Equation (8), which transforms the product of error terms into the addition of the logarithm of them, which when expanded gives us the arithmetic mean of the error given by the expression $\frac{1}{n} \sum_{i=1}^n \hat{m}_{x,i} - m_{x,i}$.

Table 1. Number Formats used for Representing LNS Numbers

Word length [†]	Int. bits	Frac. bits	Notation	Word length [†]	Int. bits	Frac. bits	Notation
5	2	2	Q(2, 2)	6	2	3	Q(2, 3)
7	3	3	Q(3, 3)	8	4	3	Q(4, 3)
9	4	4	Q(4, 4)	10	4	5	Q(4, 5)
11	4	6	Q(4, 6)	12	4	7	Q(4, 7)
13	4	8	Q(4, 8)	14	4	9	Q(4, 9)
15	4	10	Q(4, 10)	16	5	10	Q(5, 10)

[†] 1-bit reserved for sign.

This implies that calculating the geometric mean in the real domain corresponds to calculating arithmetic mean of the absolute error of the exponents, which corresponds to rounding in the log domain. Furthermore, rounding in the log domain guarantees that in non-exceptional cases the absolute rounding error is bound by half of one ULP_{LNS} . However, a problem arises when rounding to real-domain zero, which has a log domain value roughly equal to $-\infty$. Thus, when rounding a non-zero value to zero, the log-domain error is infinite. If our rounding goal is to minimize the log domain error, then we would never round a value to real domain zero. However, never rounding to zero does not make sense, at least to us. Therefore, we round values of less than half of one log domain ULP_{LNS} down to zero.

Rounding in the real domain eliminates the rounding-to-zero anomaly of an infinite error. But rounding to zero causes the real-domain relative errors to be 100%. In all other cases of rounding, the real domain relative error is bound by $\sqrt{r} - 1$ for all $r > 1$. Despite a very extensive literature on LNS over decades, we have not been able to locate a similar discussion of log-domain versus real-domain rounding. In the remainder of the article, we round in the log domain, which allows us to present errors in ULP_{LNS} (which equals 2^{-f} in the log domain), but rounding in the real domain gives similar results. Furthermore, when presenting results on the average absolute error in the log domain, we report values that underflow to zero—and thus result in an infinite log-domain error—separately from values that do not underflow.

4 CONVERSION OF NON-LNS INPUTS

When computing with LNS, it is common that the original inputs to the program are in a different format, such as **floating point (FP)**. We normally convert inputs to LNS once, and perform all subsequent computation in LNS. Any LNS and FP format contain different sets of real values. Thus, when we convert input values to LNS, we normally round each input to the nearest LNS number. Note that the rounding error for converting inputs tend to be cumulatively smaller than arithmetic rounding errors, because most algorithms perform multiple arithmetic operations for each input. In this section, we show that we can reduce the conversion rounding error, as compared to base-2 LNS, by selecting another base that better captures the distribution of the input data.

We consider a simple case of input data conversion, where the input consists of a short length FP input type, which we convert to a logarithmic number format with the same number of bits. Each FP number has a sign bit, i exponent bits, and f mantissa bits, and we convert to an LNS type with a sign bit, and an exponent with i integer bits and f fractional bits. Very-low-precision types have been used for both FP and LNS, even as low as 3-bits [25, 30]. In light of these contributions it is important to analyze the effect of quantization for shorter word lengths. The word length formats used for our work are given in Table 1.

Fig. 1. Bases with minimum quantization error for various word lengths for (a) scaled and (b) unscaled LNS data.

In our experiments, we assume that all FP numbers of the appropriate input type are equally likely to appear. In addition to considering base-2 LNS, we consider around 587 evenly distributed bases between $\sqrt{2}$ and 2. When computing the rounding error in conversion, we use the arithmetic mean of the absolute error of the exponents (in the log domain) as discussed in Section 3.2 and given in Equations (3) and (9).

The problem of rounding a non-zero FP number to zero during conversion is the same as stated earlier in Section 3.2, and we exclude values that underflow or overflow during rounding, as also stated earlier in Section 3.2. Figure 1(b) shows that for the distribution of FP numbers, the LNS base that gives the lowest quantization error is typically close to 2 for all word lengths given in Table 1.

Note, however, that each LNS base gives both a different *range* of real-domain values, and a different *distribution* of those values, with bases close to base-2 having a range closest to a floating-point distribution for a given word length. This is reflected in Figure 1(b) where bases close to base-2 provide the least quantization error across all word lengths. We consider a second experiment where we apply a constant scaling factor to all LNS values, so that the range of LNS values is the same for all bases. With this scaling, different bases lead to different distributions of LNS values, but the range is constant. We use a scaling factor that is the ratio of the maximum floating-point and LNS number for each base. Figure 1(a) shows the effect of keeping the range of LNS values constant through scaling with different bases resulting in distributions that can reduce the quantization error, compared to base-2. For very short word lengths of 5–7, the optimal base for quantization error is much lower than base-2.

The average absolute representation error, as a percentage of ULP_{LNS} in the log domain and for a scaled LNS distribution, for a few select bases is shown in Table 2 along with the difference of other bases to the best base that gives the minimum representation error, in columns two and three. Other columns describe another type of error, which will be explained shortly.

It can be seen from Table 2 that bases other than base-2 can give significant reduction in error, with a minimum 4× improvement for $Q(2, 2)$ over base-2 and a maximum of close to three order of magnitude for larger word lengths. However, the distribution and range of inputs may be domain dependent. A base that results in a distribution of values that is suitable for one domain might

Table 2. Percentage Average Relative Error for Various Bases

Number format	Base	Avg. abs. repr. error (%) [†]	Avg. abs. arith. error (%)			
			$\phi^+(x)$	$N_z^\star$	$\phi^-(x)$	$N_z^\star$
Q(2, 2)	1.414	2.6888	22.9985	0	20.5680	4
	1.741	5.4437	22.1167	0	21.5063	1
	1.417	2.8546	22.7008	0	17.1544	4
	2.0	9.7655	26.0090	0	25.4259	1
Q(2, 3)	1.442	1.2682	24.2746	0	26.2439	7
	1.999	21.6569	22.9960	0	21.1268	7
	1.415	1.6456	24.8867	0	19.9971	1
	2.0	23.6369	23.6633	0	21.4604	1
Q(3, 3)	1.659	1.6902	25.4686	0	24.7396	1
	1.496	9.6850	23.4968	0	23.1648	1
	1.802	3.5171	25.0584	0	22.1066	1
	2.0	23.6369	25.3094	0	23.7062	1
Q(4, 3)	1.851	1.6070	25.2744	0	24.6680	1
	1.730	5.0651	24.8420	0	25.1975	1
	1.802	2.4998	25.0584	0	22.1066	1
	2.0	23.6369	25.3094	0	23.7062	1
Q(4, 4)	1.850	1.2107	25.2267	0	25.6736	1
	1.697	6.1957	25.0533	0	24.0765	1
	1.973	6.5001	25.2602	0	23.4122	1
	2.0	14.9733	25.5656	0	24.7483	1
Q(4, 5)	1.852	0.6152	25.2049	0	24.6621	1
	1.718	4.5963	25.1177	0	24.3975	1
	1.666	8.1964	25.2180	0	24.0779	1
	2.0	22.1417	25.4424	0	24.9559	1
Q(4, 6)	1.831	0.3366	25.2581	0	24.8755	1
	1.414	31.6263	24.2758	0	24.2526	1
	1.422	30.7519	24.3959	0	23.9850	1
	2.0	17.6047	25.3424	0	25.1857	1
Q(4, 7)	1.836	0.1665	25.2507	0	25.0542	1
	1.467	25.9829	24.3400	0	24.2350	1
	1.460	26.7079	24.3655	0	24.2174	1
	2.0	20.8774	25.2707	0	25.2689	1
Q(4, 8)	1.836	0.0844	25.2269	0	25.0775	1
	1.522	20.4675	24.4026	0	24.1908	1
	2.0	20.3949	25.2519	0	25.2763	1
Q(4, 9)	1.834	0.0427	25.2008	0	25.2364	1
	1.579	15.1358	24.4734	0	24.4326	1
	1.586	14.5189	24.4967	0	24.3751	1
	2.0	19.9790	25.2250	0	25.2300	1
Q(4, 10)	1.834	0.0214	25.1843	0	25.1848	1
	1.641	9.8505	24.4981	0	24.4702	1
	1.643	9.6932	24.5073	0	24.4485	1
	2.0	19.6907	25.1871	0	25.1280	1
Q(5, 10)	1.915	0.0228	25.1996	0	25.1956	1
	1.445	37.2458	25.1653	0	25.1708	1
	1.503	31.3220	25.1681	0	25.0628	1
	2.0	19.6955	25.1871	0	25.1280	1

[†]Scaling factor, $S = n \implies S = \frac{\max(\text{fp})}{\max(\text{lns})}$. $\star N_z$ = Number of values rounded to zero in the real domain.

Fig. 2. The $\Phi(x)$ function for performing LNS addition and subtraction.

be a poor match for another domain. Thus, matching the LNS base to an input distribution may be useful when designing a hardware accelerator that targets a specific domain. But for general-purpose LNS processing, it is probably not worthwhile to build a base that is specific to one domain into the hardware LNS units. Our experiments show that we can adapt the LNS base to reduce the rounding when converting *some* input distribution to LNS. However, any reduction in conversion error may also be domain-dependent. In contrast, our approach in the next section improves the accuracy of arithmetic, and is less domain-dependent.

5 LNS ADDITION AND SUBTRACTION

The set of LNS numbers is not closed under addition and subtraction. For example, if the exponent consists only of integer bits (i.e., $f = 0$), then base-2 has real-domain numbers $\{1, 2, 4, 8, 16 \dots\}$. The sum of two numbers in this set may not be a member of the set. For example, $1 + 2 = 3$, but 3 is not an element of the set. In practical LNSs, we round the results of addition and subtraction to the nearest representable LNS number resulting in an inherent rounding error in the $+$ and $-$ operation regardless of their implementation.

However, different bases lead to the LNS containing different sets of values. For example, the set of LNS numbers with an integer exponent and base- $\sqrt{2}$ contains numbers of the form

$$\{1, \sqrt{2}, 2, 2\sqrt{2}, 4, \dots\}.$$

In this number system, if we add $1 + 2 = 3$, then the nearest representable LNS number is $2\sqrt{2} \approx 2.828$. Thus, both the set of representable values and the size of rounding errors depends on the base. A common way to compute LNS addition and subtraction uses the identities [20]:

$$m_{\text{add/sub}} = \max(m_x, m_y) + \Phi(x), \quad (10)$$

where

$$\Phi(x) = \begin{cases} \Phi^+(x) = \log_b |1 + b^{-x}| & S_\Phi = 0 \\ \Phi^-(x) = \log_b |1 - b^{-x}| & S_\Phi = 1 \end{cases}, \quad (11)$$

where b is the base, $S_\Phi = s_x \oplus s_y \oplus op$ (op is 0 for addition and 1 for subtraction) and $x = |m_x - m_y|$ (where $x \leq 0$). The functions $\Phi^+(x)$ and $\Phi^-(x)$ are non-linear, as shown in Figure 2

The challenge of computing function $\Phi(x)$ can offset the gains achieved by other useful properties of LNS. A number of methods exist such look-up tables, CORDIC, Taylor series expansion or function approximation [20, 26]. We use the look-up table method, which requires $O(2^{i+f})$ space complexity, and is suitable for word lengths up to 20-bits [41]. For the low-precision word sizes we consider, in the range 5–16 bits, the look-up tables are relatively small.

We represent the Φ^+ and Φ^- functions as look-up tables for different bases. LNS with fixed-point exponent is not closed under addition or subtraction, so we round the values in the tables to the nearest representable LNS number. Different LNS bases lead to number systems containing different sets of numbers. Thus, these look-up tables are different for each base, and therefore the rounding errors are different for each base.

Typically, in an LNS arithmetic unit, look-up tables are implemented as simple **read-only memories (ROMs)**. However, in some cases, the table entries for the two functions are trivial with a very few non-zero values and thus realizing them as logic gates can be beneficial than implementing them as ROMs.

5.1 Realization of Large LNS/Small Linear Values in the Subtraction Tables

According to Equation (11), the first entry in the tables correspond to when b^{m_x} and b^{m_y} are equal. This entails two things. First, output of the subtraction of such numbers will be zero and that for $\phi^-(x)$, the first value corresponds to $\log_b |0|$, which is $-\infty$. Similar to this, for the next few values, the difference between b^{m_x} and b^{m_y} is very small implying that subtraction between these two values results in a very small number. This also means that $\log_b |1 - b^{-x}|$ of Equation (10) returns a very large negative value.

These large negative values in the subtraction table correspond to very small values in the real domain, which are smaller than the smallest representable number. This number corresponds to $b^{-(2^i - 2^{-f})}$. One needs to make a decision about their representation in the tables. In the real domain, values that are smaller than the smallest representable number are either rounded to zero or the smallest representable number, depending on the distance between the number itself and the two adjacent numbers. However, one cannot represent real zero in the LNS domain as that is equivalent to $-\infty$. Thus, the underlying question is the representation of linear zero/LNS $-\infty$ in the subtraction tables.

Since in these cases the output of subtraction either results in a zero or a very small value that is to be rounded to zero in the real domain, we can utilize this by storing zeros in the table and having a small logic circuit that can detect such cases and forcing the output to zero. Storing zeros simplifies the logic circuitry synthesized for realizing the subtraction tables.

5.2 Evaluation of Arithmetic Error

We computed the arithmetic error as a percentage of ULP_{LNS} in the log domain (where $ULP_{LNS} = 2^{-f}$) and considered bases ranged from $\sqrt{2}$ and 2, similar to what was done for representation error. The rounding was computed in the log domain with the error measured as defined in Section 3.2.

The result is shown in Figures 3(a) and 3(b) for various word lengths ranging from 5 to 16 bits (including one sign bit) where only the base that achieves the minimum arithmetic error for a given word length is shown. The error measures do not include errors generated when values are rounded to zero in the real domain and remain within the bounds given by Equation (7). Also shown in Figures 3(c) and 3(d) is the variation in the average absolute arithmetic error as a percentage of ULP_{LNS} (in the log domain) for $Q(4, 3)$ and that for this particular case, base-1.730 and base-1.802 provide the lowest arithmetic error. This is also indicated in Table 2, where the results for arithmetic error were presented along with representation error. This is a 2% and 7% improvement for the addition and subtraction tables over standard base-2.0.

When rounding to nearest, the average arithmetic error will be around 25% of the ULP_{LNS} with the maximum absolute error being $\leq 0.5 \times ULP_{LNS}$. Figures 3(a) and 3(b) show that base-2 is not the best choice, especially for smaller word lengths.

As mentioned in Section 4, the errors generated due to representing an input distribution is highly domain dependent. However, the values in addition/subtraction tables are dependent on the particular arithmetic operation and the base in which they are represented and this result shows that rounding error for LNS arithmetic can be reduced by selection of an appropriate base.

However, different bases give minimum rounding error for the addition and subtraction tables for different word lengths. These bases do not correspond to the bases that give minimum

Fig. 3. Bases with minimum arithmetic error for various word lengths for (a) addition and (b) subtraction tables. Average absolute arithmetic error for $Q(4, 3)$ as a percentage of one ULP_{LNS} (in the log domain) for (c) addition and (d) subtraction tables.

representation error as shown in Figure 1. To compare the arithmetic and representation error for different interesting bases, Table 2 shows them for each of these bases. It also identifies the best base for each type of error. Also shown are the number of values that are rounded to zero in the real domain. These values are not counted when computing the arithmetic error, for reasons explained in Section 3.2.

Typically, rounding the fixed-point exponents to nearest results in a maximum error of 0.5 ULP_{LNS} and average error of 0.25 ULP_{LNS} . However, it can be seen in Table 2 that for very short word lengths, the average arithmetic error for bases other than base-2 is much lower than $0.25 \times ULP_{LNS}$, especially for the subtraction table. For example, base-1.417 gives an arithmetic error of 17.1544% of ULP_{LNS} for $Q(2, 2)$ while base 1.415 gives 19.9971% of ULP_{LNS} (in the log domain) for $Q(2, 3)$. These errors converge to around $0.25 \times ULP_{LNS}$ as we increase the word length.

Furthermore, as compared to base-2, significant improvements can be achieved as compared to base-2 of up to 17.6% and 48.2% for addition and subtraction tables, respectively, for $Q(2, 2)$. It can also be seen that the representation error for bases that give minimum arithmetic error is typically within the 25% range expected of average rounding error, bar only one exception when the word length specification matches that of half-precision floating point ($Q(5, 10)$).

6 SYNTHESIS OF ADDITION AND SUBTRACTION TABLES

A number of methods exist in literature for realizing Equation (10) in hardware, as mentioned in Section 5. However, since we are dealing with short word lengths, table-based approach is the preferred choice [41].

The functions of $\Phi^\pm(x)$ are evaluated by storing pre-computed values in a look-up table and then reading the appropriate value based on the input. The values in the table are quantized to the given word length, various choices for which are listed in Table 1. Typically, look-up tables are realized using read-only ROMs. However, if the table entries are fairly simple, they can be realized using logic gates rather than ROMs. One can also realize some columns of the table using logic gates and others using a ROM depending on the complexity of the columns.

Table 3. Implementation Cost of Various Options for a 7×128 ROM-based Table Implementation in Terms of Number of Transistors

Type of ROM	Storage	Precharge	Sense amplifier	Decoder	Total
ROM with one decoder	896	7	35	1,022	1,960
ROM with two decoders	896	56	280	$122 + 54 = 176$	1,408
5-bit ROM with two decoders	640	40	200	$122 + 54 = 176$	1,056

Thus, to realize tables, one has three options:

- Synthesize the whole table as a ROM
- Synthesize the whole table using logic gates
- Synthesize parts of the table using logic gates and parts using a ROM

A simple, fixed ROM usually consume only one transistor per bit for data storage [13]. However, even a simple fixed ROM will have additional circuitry as a row/column decoder, sense amplifiers and precharge transistors [45]. A simple differential sense amplifier will consist of five transistors [35], while one needs one precharge transistor per one column of bits. The address decoding for a memory is typically performed by a combination of row and column decoders if memories are arranged as a matrix of words rather than a one dimensional array of individual words or only one address decoder if arranged as a one dimensional array structure.

In the case of realizing the tables using logic gates, one can either synthesize the logic for individual column of bits independently or combine the columns to enable the synthesis tool to share logic to realize various columns of bits. In our experiments, the total synthesis cost, which we present in terms of transistors, for a combined synthesis of columns of bits is lower than synthesizing individual columns separately. We use the open source ABC synthesis tool [27] and provide the tool with tables described using the **Berkeley logic interchange format (BLIF)** file [7] and a library containing only NAND, NOR and Inverter gates to enable us to estimate the number of transistors. Each NAND and NOR gate is typically made up of four transistors while the inverter is made up of two transistors [35]. This can be described as a first degree estimation, as we do not take into account the differences in the sizes of the two different types of transistors that are used to realize the gates using the complementary MOS technology [35].

However, the question about when to use logic gates and when to use ROM will depend on a threshold. The threshold will be the point when the cost of implementing using gates goes beyond that of ROM, which has a constant cost of implementation.

Consider the example of 8-bit LNS, which requires only 128×7 with the sign bit not stored. If realized as an array structure, then the cost for such a ROM will be 896 transistors for the storage, 7 transistors for the precharge, $5 \times 7 = 35$ transistors for sense amplifiers and 1,022 transistors for a 7×128 decoder. The implementation cost for the decoder was estimated by synthesizing the decoder using the ABC synthesis tool, which was provided a BLIF file describing the truth table of it. This makes for a total cost of 1,960 transistors.

However, if we implement the memory as a matrix of 16×8 words, then one will need a 4×16 row and 3×8 column decoder, which cost 122 and 54 transistors, respectively, making the total cost of the memory to be: $896 + 8 \times 7(\text{precharge}) + 8 \times 7 \times 5(\text{sense amplifier}) + 122 + 54 = 1,408$ transistors. This number can be further reduced by investigating the values each of the addition and subtraction table takes. For the subtraction table, one needs all 7 bits of output, however, for the addition table, one can discard the two most significant bits (as they are zero) and only store the least significant five bits. This is also evident by Figure 2 where the range of $\phi^+(x)$ is limited between zero and one while that of $\phi^-(x)$ is much larger. This can reduce the total cost for the addition table to 1,056 transistors. All of this is summarized in Table 3.

Fig. 4. Number of transistors required to realize 8-bit (a) addition and (b) subtraction tables for all bases between $\sqrt{2}$ and 2.

Fig. 5. Number of transistors required for realizing successively increasing number of columns for 8-bit (a) addition and (b) subtraction tables for bases with minimum ($\text{base}_{\phi^+} : 1.832, \text{base}_{\phi^-} : 1.913$) and maximum transistor count ($\text{base}_{\phi^+} : 1.423, \text{base}_{\phi^-} : 1.424$).

Comparing the cost of ROM with logic gates for realizing complete table shows that for all bases, tables realized using gates has lower area as compared to ROM, as shown in Figure 4 for 8-bit LNS for all combination of bases between $\sqrt{2}$ and 2, with a minimum savings of 55% and 43% for the addition and subtraction tables, respectively.

6.1 Tables using Gates and ROM

Another approach to realize these tables is to realize some columns using logic gates and some using ROM. This may appear beneficial, because for those columns with only a few 1s, a gate-based realization will result in very small circuits while for the columns where changes in the value are too frequent, a ROM-based implementation may be better or of similar complexity as compared to a gate-based implementation. However, experimental evaluation shows that is not the case, as shown in Figure 5 for select bases that give the minimum and maximum transistor count for each table. The x-axis in each sub-figure shows the number of memory columns where **rev.** means that columns are combined starting from least significant column while the other alternative shows the starting point is the most significant columns. For the addition table, the two most significant columns are not shown as they are all zeros.

Fig. 6. Delay associated with logic circuit required to realize 8-bit (a) addition and (b) subtraction tables for all bases between $\sqrt{2}$ and 2.

Figure 5 shows that the maximum transistor count for the addition table increases from 26 transistors for a single column to 476 transistors for all columns combined using logic gates for base-1.423. This is an improvement of 55% to 93%. Similarly, the savings in transistor count for the subtraction case, in the worst case, is 43% to 91%. Even when combining the columns starting from the most significant column results in lower transistor count as compared to realizing the tables using a ROM.

Thus, considering only the perspective of transistor count, realizing the tables using only logic gates results in a lower transistor count as compared to a simple ROM. Figure 5 also indicates that for 8-bit LNS, a base of 1.832 and 1.913 gives minimum transistor count for addition and subtraction table, respectively. Thus, with a selection of base other than base-2, one can also reduce the implementation cost.

6.2 Delay of Logic Circuit for Tables

Typically, for a given word length, the latency of reading a word from a ROM will be the same across all reads, specially if the word being read are in different rows. In a typical implementation, the row address is asserted first to read the whole row (consisting of multiple words) into the sense amplifiers from where the exact word is selected using the column address. The latency of the two decoders (4×16 and 3×8 , as given in Section 6) each for 7- and 5-bit memories is 4.7 and 3.3 time units, respectively. Only considering the sum of the delay of the two decoders and ignoring the delay associated with the actual reading of the value from transistors, sense amplifiers and any refresh times required, the total delay is 8 time units.

However, the delays associated with realizing the tables as logic gates for different bases and 8-bit output is shown in Figure 6. The difference between the delay associated with logic gates and ROM-based implementation is not significant. With a very significant difference between the transistors required to realize the two alternatives for realizing the addition and subtraction tables, as shown in Figures 4 and 5, it can be concluded that it is better to use logic gates for realizing the tables for 8-bit LNS.

7 DECIDING THE TRADE-OFF FOR ADDITION AND SUBTRACTION TABLES

Four parameters have been discussed in the previous sections:

- Representation error
- Arithmetic error for $\phi^+(x)$ and $\phi^-(x)$
- Area in terms of transistors count for realizing tables using gates for $\phi^+(x)$ and $\phi^-(x)$
- Delay for realizing tables using gates for $\phi^+(x)$ and $\phi^-(x)$

Fig. 7. Radar plot showing the trade-offs between various parameters for $Q(3, 3)$.

For each parameter, there are unique bases in the LNS that gives the corresponding minimum value. The most important of these parameters are the two types of errors.

The trade-offs involved for one particular word length of $Q(3, 3)$ is shown in Figure 7 for various bases. Similar trade-offs exist for other word lengths. All parameters have their importance in the overall design and implementation and each base is optimal in one of the four parameters defined above but none of the base is optimal in all of the four parameters. This figure emphasizes our point of the need of analyzing different bases and the priorities while designing and implementing any system. However, it has been clearly shown that base-2.0 is not optimal under different word lengths and for different parameters.

8 HIGHER-PRECISION TABLES FOR VERY SHORT WORD LENGTHS

As mentioned earlier, a number of works use very short word lengths for realizing different systems for deployment on edge devices. However, short word lengths can introduce significant precision errors for addition and subtraction. One way to circumvent this is to use tables with higher precision for table values (referred from here on as mixed-precision tables for exploratory convenience) and rounding the result only after all required operations for, say, computing the sum of products of data and coefficients in an FIR filter are complete. Thus, there is only one source of arithmetic error at the output (at the cost of higher hardware cost).

The number format given in Table 1 was used for the analysis with fractional word length successively increased up to a maximum of total word length of 16 (including the sign bit) for each case. For example, for the case of $Q(2, 2)$, the input word length is 5-bits while the output word length is varied from 5 to 16 bits. The results for two different total input word length, $Q(2, 2)$ and $Q(3, 3)$, are shown in Figure 8, where each value represents the minimum percentage average absolute arithmetic error. For each value, the base that gives this minimum value is unique. As discussed earlier in Section 5.2, for shorter word lengths, the arithmetic error can be reduced much lower than 25% of ULP_{LNS} for bases other than base-2.

Fig. 8. Best case average absolute arithmetic error as a percentage of the ULP when the addition and subtraction mixed-precision tables are quantized to increasing fractional word length for (a) Q(2, 2) and (b) Q(3, 3) number formats identified in Table 1.

Fig. 9. Bases with minimum arithmetic error for various word lengths for (a) addition and (b) subtraction mixed-precision tables.

For each number format given in Table 1 and all output word lengths, there is no single optimal base that gives the least arithmetic error for various combinations of integer and fractional bits. This phenomenon is further emphasized in Figure 9, which identifies the various bases that give the least arithmetic error for total output word length. The legend shows the original total word length (without extended fractional bits). For example, the plot shows the optimal base for output word lengths ranging from 5 to 16 bits for the case of Q(2, 2). Again it re-emphasizes the earlier reached conclusion that bases other than base-2 should be considered when designing a system as it can give lower error and implementation cost.

8.1 Synthesis of Higher-precision Tables for Very Short Word Lengths

The other key issue to deal with extra fractional bits is the size of these tables, once they are realized in hardware. As shown in Figures 4 and 5, it is better to realize the tables using gates. However, with increasing output word length for the tables, one needs to analyze whether the conclusion reached in Section 6 still stands. For this, the addition and subtraction tables were synthesized using both ROM and standard cell gates for all input and output word length combinations. Table 4 shows the number of transistors required to realize the ROMs based on the model earlier defined in Section 6. The size of decoders for the various tables were chosen to minimize the transistor count.

Table 4. Implementation Cost for ROM-based Table Implementation, for Mixed-precision Tables, in Terms of Number of Transistors

Input word length (WL_i)	Transistor count for ROM (with various output word lengths (WL_o))											
	5	6	7	8	9	10	11	12	13	14	15	16
5	168	196	224	252	280	308	336	364	392	420	448	476
6	–	345	401	457	513	569	625	681	737	793	849	905
7	–	–	661	749	837	925	1,013	1,101	1,189	1,277	1,365	1,453
8	–	–	–	1,408	1,584	1,760	1,936	2,112	2,288	2,464	2,640	2,816
9	–	–	–	–	2,736	3,040	3,344	3,648	3,952	4,256	4,560	4,864
10	–	–	–	–	–	5,844	6,452	7,060	7,668	8,276	8,884	9,492
11	–	–	–	–	–	–	11,832	12,952	14,072	15,192	16,312	17,432
12	–	–	–	–	–	–	–	25,400	27,640	29,880	32,120	34,360
13	–	–	–	–	–	–	–	–	52,728	57,016	61,304	65,592
14	–	–	–	–	–	–	–	–	–	113,020	121,596	130,172
15	–	–	–	–	–	–	–	–	–	–	237,312	254,080
16	–	–	–	–	–	–	–	–	–	–	–	506,112

Fig. 10. Bases with minimum transistor count for various word lengths for the (a) addition and (b) subtraction mixed-precision tables.

The addition ($\phi^+(x)$) and subtraction ($\phi^-(x)$) tables were also synthesized for all input and output word length combinations using standard cell logic gates. For each case, the base that gives the minimum transistor count is unique and a distribution of these bases is shown in Figure 10. This figure, together with the results shown in Figure 9, identifies the various bases that give minimum arithmetic error and implementation cost, respectively, with lower the word length, the more smaller bases give minimum transistor count.

However, arithmetic error is more important than transistor count to ensure correctness of result. Although the hardware cost of realizing tables with bases giving minimum arithmetic error is larger than the best case in terms of transistor count, it is still significantly better than realizing it in ROMs. Results for the subtraction table are shown in Table 5 as typically it is more complex to realize than the addition table and shows an order of magnitude improvement, with a minimum reduction of 168% in transistor count to a maximum of nearly 10× reduction.

It can also be seen from Table 5 that as the output word length increases, for each input word length, the transistor count does not increase significantly, making the mixed-precision approach an attractive one given the reduced error measurement.

Table 5. Implementation Cost for Gates-based Table Implementation, for Mixed-precision Table for $\phi^-(x)$, in Terms of Number of Transistors for Bases that give the Minimum Arithmetic Error

Input word length (WL_i)	Transistor Count for ROM (with various output word lengths (WL_o))											
	5	6	7	8	9	10	11	12	13	14	15	16
5	64	96	122	120	156	186	208	214	244	254	282	290
6	–	158	244	264	394	416	486	452	530	614	662	788
7	–	–	456	454	686	808	834	810	1,114	1,026	1,336	1,484
8	–	–	–	494	530	608	872	1,218	1,074	1,768	1,428	2,582
9	–	–	–	–	878	1,326	1,744	2,996	2,122	3,962	3,018	4,460
10	–	–	–	–	–	1,746	3,604	4,196	5,150	6,040	4,310	7,528
11	–	–	–	–	–	–	6,132	7,130	8,916	10,080	10,664	12,458
12	–	–	–	–	–	–	–	10,448	13,072	15,884	16,184	19,132
13	–	–	–	–	–	–	–	–	18,962	21,918	27,644	31,194
14	–	–	–	–	–	–	–	–	–	33,438	41,068	50,580
15	–	–	–	–	–	–	–	–	–	–	63,140	75,726
16	–	–	–	–	–	–	–	–	–	–	–	66,608

9 LNS ADDER AND MULTIPLIER DESIGN IN FIXED- AND FLOATING-POINT NUMBER SYSTEM

As identified in Sections 1 and 2, LNS finds application in signal processing and neural networks. Both of these systems are based heavily on the operation of addition and multiplication. The key advantage of LNS over fixed- and floating-point stem from its inherent simplification of multiplication into addition. However, the LNS can only approximate the addition operation, since addition is not closed under LNS, as described in Section 5.

Multiplication in LNS is equivalent to adding the exponents. Thus, the circuit to realize an LNS multiplier is similar in complexity to that of a fixed-point adder [23]. To ascertain which number system to use from a hardware complexity purpose it is important to realize the hardware circuit for LNS addition/subtraction and compare it against fixed-point multiplier. Since in applications like neural networks the dominant number system used is floating point, it is important to realize floating-point addition and multiplication circuits to analyze which number system provides the smallest circuit.

The circuit to perform addition and subtraction in LNS is given in Figure 11 and is a modified form of the circuit given by Kenny et al. [20]. In the circuit, the functions of $\phi^+(x)$ and $\phi^-(x)$ are implemented using look-up tables, as mentioned in Section 6. The zero detect module shown takes into account whether the two numbers being added are both zero (shown by the z_x bit of Equation (1)), the difference of two number is zero or so small that the result will be rounded to zero in the real/linear domain, as discussed in Section 5.1.

Addition and multiplier circuits were synthesized for logarithmic, fixed- and floating-point number systems for word length ranging from five to 16-bits using Cadence Genus with the 45-nm GSCl liberty cell library from the Free45PDK. For LNS, the bases were varied from $\sqrt{2}$ to 2. The results are shown in Table 6, where $Q(i, f)$ indicates i integer and f fractional bits and does not include the sign bit. For fixed point, it needs to be read as $Q(i + 1, f)$, since 2's complement number representation is assumed. Both adder (Float.A) and multiplier (Float.M) floating-point results are presented.

As evident in Table 6, the LNS adder requires lower or similar area as compared to fixed-point multipliers for short word lengths up to $Q(4, 4)$. Although the floating-point multiplier is the smaller of all circuits (for most of the cases), the LNS multiplier and fixed-point adder are much

Fig. 11. Circuit schematic for LNS addition and subtraction.

Fig. 12. Radar plot showing the trade-offs between the three design parameters for $Q(3, 3)$.

simpler and smaller. Thus, overall, floating point will not have the smallest of the circuit. Also, for LNS the same circuit is used to perform both addition and subtraction, which opens up the possibility of re-using the same circuit for the two operations if the input sample rate allows for some level of multiplexing. Furthermore, for LNS, the adder and subtractor were realized using the look-up table approach for implementing the functions $\phi^+(x)$ and $\phi^-(x)$ in Equation (10). They can also be realized using different techniques like weighted sum of bit-products [20], which can reduce the area complexity. However, in this work all error calculations are based on table-based implementations, the same method was used for realizing circuits for consistency.

The results in Table 6 for each number format shows the best LNS base other than base-2 giving the lowest area. This does not correspond to the shortest delay through the circuit or the lowest power consumption. Figure 12 gives a better insight into the inherent trade-offs involved here.

10 FINITE-IMPULSE RESPONSE FILTER: AN APPLICATION

Frequency selective FIR filters are used in many applications and are therefore a key building block in a digital signal processing system [29]. The difference equation that defines the output of

Table 6. A Comparison of Synthesis Results for Fixed-point Multiplier, Floating-point Adder and Multiplier, and LNS Adder Circuits

Number format	Circuit type	Area (μm^2)	Delay (ps)	Power (μW)	Number format	Circuit type	Area (μm^2)	Delay (ps)	Power (μW)
$Q(2, 2)$	Fixed	495	874	101	$Q(2, 3)$	Fixed	688	1,035	146
	Float.M	415	991	80		Float.M	576	1,467	101
	Float.A	565	1,868	108		Float.A	706	2,325	120
	LNS (2.0)	445	1,316	109		LNS (2.0)	577	1,650	148
	LNS (1.421)	423	1,214	107		LNS (1.443)	564	1,583	147
$Q(3, 3)$	Fixed	890	1,190	195	$Q(4, 3)$	Fixed	1,143	1,328	251
	Float.M	680	1,504	126		Float.M	730	1,620	135
	Float.A	913	2,922	170		Float.A	1,002	3,314	170
	LNS (2.0)	744	1,900	170		LNS (2.0)	883	2,170	197
	LNS (1.718)	713	1,832	175		LNS (1.999)	857	2,061	191
$Q(4, 4)$	Fixed	1,332	1,515	303	$Q(4, 5)$	Fixed	1,599	1,753	379
	Float.M	962	1,931	180		Float.M	1,102	2,307	216
	Float.A	1,145	3,179	222		Float.A	1,305	3,557	249
	LNS (2.0)	1,325	2,477	263		LNS (2.0)	1,622	2,477	263
	LNS (1.96)	1,244	2,660	264		LNS (1.948)	1,572	2,930	347
$Q(4, 6)$	Fixed	1,856	1,890	458	$Q(4, 7)$	Fixed	2,208	2,109	541
	Float.M	1,328	2,518	258		Float.M	1,553	2,831	318
	Float.A	1,460	4,118	288		Float.A	1,653	4,688	350
	LNS (2.0)	2,643	3,374	552		LNS (2.0)	4,319	4,342	824
	LNS (1.92)	2,523	3,316	550		LNS (1.997)	4,113	4,022	804
$Q(4, 8)$	Fixed	2,531	2,502	646	$Q(4, 9)$	Fixed	2,842	2,652	726
	Float.M	1,796	3,008	380		Float.M	2,062	3,364	438
	Float.A	1,769	4,528	355		Float.A	1,958	5,358	423
	LNS (2.0)	7,468	5,419	1269		LNS (2.0)	13,155	6,328	1,897
	LNS (1.998)	7,289	5,176	1234		LNS (1.975)	13,032	6,341	1,906

Fig. 13. Direct form FIR filter.

an N th-order (length $N + 1$) FIR filter is

$$y(n) = \sum_{k=0}^N h(k)x(n - k), \quad (12)$$

where $y(n)$ is the output sequence, $x(n)$ is the input sequence, and $h(k)$ are the coefficients. The direct form of FIR filter to realize Equation (12) is shown in Figure 13.

The arithmetic complexity of FIR filters primarily depends on the number of multiplications, which according to Equation (12) is proportional to the filter order. The filter order is, for

Table 7. A Comparison of Synthesis Results for FIR Filters Implemented using Fixed, Floating, and Logarithmic Number Systems

Number format	Circuit type	Area (μm^2)	Delay (ps)	Power (μW)	Number format	Circuit type	Area (μm^2)	Delay (ps)	Power (μW)
$Q(2, 2)$	Fixed	4,914	193	1,700	$Q(2, 3)$	Fixed	6,841	197	2,424
	Float	7,891	414	3,346		Float	10,877	411	5,200
	LNS (2.0)	3,663	185	1,542		LNS (2.0)	5,301	189	2,443
	LNS (1.414)	3,454	185	1,297		LNS (1.417)	4,620	185	1,828
$Q(3, 3)$	Fixed	8,964	202	3,239	$Q(4, 3)$	Fixed	11,456	208	4,251
	Float	12,554	458	5,694		Float	13,693	524	5,275
	LNS (2.0)	6,691	194	3,007		LNS (2.0)	8,393	195	3,231
	LNS (1.975)	6,568	194	2,957		LNS (1.96)	8,044	195	3,117
$Q(4, 4)$	Fixed	14,104	207	5,269	$Q(4, 5)$	Fixed	16,857	216	6,601
	Float	16,418	540	6,042		Float	19,250	560	7,907
	LNS (2.0)	11,945	190	4,810		LNS (2.0)	16,954	195	7,406
	LNS (1.96)	12,436	194	4,969		LNS (1.998)	16,554	195	7,223
$Q(4, 6)$	Fixed	20,860	246	8,091	$Q(4, 7)$	Fixed	23,762	244	9,373
	Float	23,351	587	9,999		Float	27,019	619	11,221
	LNS (2.0)	27,674	190	12,295		LNS (2.0)	45,148	195	18,193
	LNS (1.92)	27,183	190	12,019		LNS (1.997)	44,707	190	11,896
$Q(4, 8)$	Fixed	27,866	256	11,087	$Q(4, 9)$	Fixed	31,347	251	12,698
	Float	30,055	674	13,238		Float	37,063	659	15,873
	LNS (2.0)	81,810	189	32,375		LNS (2.0)	146,520	189	47,353
	LNS (1.998)	79,898	195	31,760		LNS (1.975)	149,660	185	48,731

linear-phase single stage FIR (SSF) filters, roughly proportional to the inverse of the width of the transition band [29].

LNS multiplication is performed by adding the fixed-point exponents, a simple operation. However, addition/subtraction in LNS is non-trivial and requires a table-based implementation as shown in Figure 11. As shown in Section 9, LNS provides lower hardware cost in terms of area for short word length adders/subtractors as compared to its counterparts and it is important to analyze whether these benefits extend to FIR filters.

For this analysis, an 11th-order FIR filter was implemented in hardware using the same setup as described in Section 9 with the three number systems. The input, output and filter coefficients were assumed to be in respective number systems. For LNS, the FIR filter was realized for a select number of bases identified earlier in Tables 2 and 6. The results of this analysis is shown in Table 7.

The results shown in Table 7 are consistent with those shown in Table 6. FIR filters realized using logarithmic number systems achieve a lower area (minimum savings of 2%) while also having the shortest delay (minimum improvement of 4%) and lower power consumption (minimum improvement of 6%) as compared to its fixed- and floating-point counterparts up to the word length corresponding to $Q(4, 4)$. Overall, LNS-based FIR filters also are the fastest even though they occupy an area that is, at a maximum, nearly $3\times$ more than others.

The large area for LNS-based FIR filters for larger word lengths is due to the exponential increase in the memory size required for implementing the $\phi^+(x)$ and $\phi^-(x)$ functions, as given in Equation (10). This can be reduced by using approaches like the co-transformation approach by Arnold et al. [5], the multi-table approach proposed by Coleman et al. [41] or the multipartite table methods proposed by Dinechin et al. [15].

Table 8. Cost of Conversion between Logarithmic and Fixed-point Number Systems in Terms of Area (μm^2)

	$Q(2, 2)$		$Q(2, 3)$		$Q(3, 3)$		$Q(4, 3)$		$Q(4, 4)$	
Base	2.0	1,414	2.0	1,417	2.0	1,975	2.0	1,96	2.0	1,96
Fixed $\rightarrow$ LNS	52	45	149	120	310	358	512	609	1,325	1,429
LNS $\rightarrow$ Fixed	29	30	68	111	146	185	308	217	608	651
	$Q(4, 5)$		$Q(4, 6)$		$Q(4, 7)$		$Q(4, 8)$		$Q(4, 9)$	
Base	2.0	1,998	2.0	1,92	2.0	1,997	2.0	1,998	2.0	1,975
Fixed $\rightarrow$ LNS	2,721	2,715	5,362	5,433	10,879	10,751	19,631	20,418	38,118	38,783
LNS $\rightarrow$ Fixed	1,413	1,453	2,609	2,782	5,353	5,265	10,195	10,347	20,003	20,577

Typically conversion can be carried out offline and all calculations can be done online in LNS. However, to show the cost of conversion between logarithmic and fixed point, conversion circuits were realized as look-up tables and synthesized using the same process as the FIR filter. Only conversion between LNS and fixed point is shown, since FIR filters are typically implemented using fixed point. The results are shown in Table 8 and it can be seen that the cost of realizing FIR filters using LNS is still lower than using fixed point, albeit with reduced margins, including conversion results, up to word length of $Q(4, 4)$. Optimized circuits for conversion have been proposed in the literature by Nam et al. [31], Zhu et al. [46], and others, which can help in reducing this cost.

11 CONCLUSION

LNSs are a valuable alternative to floating point for embedded systems and domain-specific hardware accelerators. Existing LNS research has almost exclusively used base-2. Although other bases, such as base- $\sqrt{2}$ or base- $\sqrt[3]{2}$, have been explored, we show that these are simple aliases of base-2, which can be constructed by moving the binary point of the fixed-point exponent in a base-2 LNS.

In this article, we show that non-base-2 LNSs can have significant advantages, particularly for short word lengths from 5 to 16 bits. We first define a numeric error for LNS that measure ULP_{LNS} as an absolute error in the fixed-point exponent of the LNS, which corresponds to a multiplicative error in the corresponding real domain value. We show that errors in conversion to LNS can be reduced dramatically by choosing an appropriate base and scaling factor, at least for the distribution of FP numbers. In our experiments, we reduce the typical conversion error from 0.20 to 0.25 ULP_{LNS} to less than 0.02 ULP_{LNS} .

We note that LNSs with a fixed word size are not closed under addition or subtraction, which leads to rounding errors in these operations. We also show that for a given LNS word size, some LNS bases lead to inherently lower average rounding errors. When rounding to nearest, the maximum arithmetic error for non-special values is 0.5 ULP_{LNS} , and we expect the average to be around 0.25 ULP_{LNS} . For example, in 5-bit (i.e., $Q(2,2)$) LNS, we show that base-2 leads to average errors of 0.260 (add) and 0.254 (subtract) ULP_{LNS} . In contrast, base-1.417 gives an average error of just 0.227(add) and 0.172 (subtract) ULP_{LNS} . Over millions of operations, even small differences in rounding errors can accumulate to large errors.

Although LNS add/subtract are often implemented with low-latency ROM lookup tables, we show that for small word sizes implementing these functions in dedicated logic is also efficient. The choice of base affects the truth table of the functions computed for LNS add/subtract, and can therefore affect the area and latency of the add/subtract units. By appropriate choice of base, the logic area of the $\phi^+(x)$ and $\phi^-(x)$ logic can be reduced by around 2 $\times$. The building blocks of adder, subtractor and multiplier are implemented using logarithmic, fixed- and floating-point numbers

and extended to a complete FIR filters, and we showed that LNS can be beneficial from all aspects of area, delay and power for very short word lengths.

We conclude that the choice of base affects the conversion errors, arithmetic errors, and hardware implementation of arithmetic units on several dimensions and that LNS can be used to implement different systems, like the FIR filter. Different bases allow different trade-offs between these goals. Although prior research in LNS has generally assumed base-2, our results show that other bases can reduce arithmetic and conversion errors while simultaneously reducing the area and latency of add/subtract and FIR units. Thus, when designing a low-precision LNS for custom hardware, one should consider looking beyond base-2.

ACKNOWLEDGMENTS

We extend our gratitude to Dr. Oscar Gustafsson of Linköping University, to the Institute of Technology Carlow, and to the anonymous reviewers for helpful feedback and comments.

REFERENCES

- [1] S. A. Alam and D. Gregg. 2020. Beyond base-2 logarithmic number systems (WiP Paper). In *Proceedings of the ACM SIGPLAN/SIGBED International Conference on Language, Compilers, and Tools for Embedded Systems*. ACM.
- [2] S. A. Alam and O. Gustafsson. 2014. Design of finite word length linear-phase FIR filters in the logarithmic number system domain. *VLSI Des.* 2014 (2014), 14.
- [3] M. Arnold, E. Chester, and C. Johnson. 2020. Training neural nets using only an approximate tableless LNS ALU. In *Proceedings of the IEEE International Application-Specific System Architecture Processors Conference*. 69–72.
- [4] M. G. Arnold. 2002. Reduced power consumption for MPEG decoding with LNS. In *Proceedings of the IEEE International Application-Specific System Architecture Processors Conference*. 65–67.
- [5] M. G. Arnold, T. A. Bailey, J. R. Cowles, and M. D. Winkel. 1998. Arithmetic co-transformations in the real and complex logarithmic number systems. *IEEE Trans. Comput.* 47, 7 (1998), 777–786.
- [6] C. Basetas, I. Kouretas, and V. Paliouras. 2007. Low-power digital filtering based on the logarithmic number system. In *Proceedings of the International Workshop on Power Timing Modeling Optimization Simulation*. Vol. 4644. 546–555.
- [7] BLIF 1992. Berkeley logic interchange format (BLIF). University of California, Berkeley. Retrieved from <http://www.cs.columbia.edu/~cs6861/sis/blif/index.html>.
- [8] N. Bruschi, A. Garofalo, F. Conti, G. Tagliavini, and D. Rossi. 2020. Enabling mixed-precision quantized neural networks in extreme-edge devices. In *Proceedings of the ACM International Conference on Computing Frontiers*. 217–220.
- [9] A. Capotondi, M. Rusci, M. Fariselli, and L. Benini. 2020. CMix-NN: Mixed low-precision CNN library for memory-constrained edge devices. *IEEE Trans. Circ. Syst. II: Express Briefs* 67, 5 (2020), 871–875.
- [10] D. V. Satish Chandra. 1998. Error analysis of FIR filters implemented using logarithmic arithmetic. *IEEE Trans. Circ. Syst. II* 45, 6 (June 1998), 744–747.
- [11] D. V. S. Chandra. 1998. Error analysis of FIR filters implemented using logarithmic arithmetic. *IEEE Trans. Circ. Syst. II* 45, 6 (1998), 744–747.
- [12] J. Chen and X. Ran. 2019. Deep learning with edge computing: A review. *Proc. IEEE* 107 (2019), 1655–1674. Issue 8.
- [13] Wai-Kai Chen (Ed.). 2003. *Memory, Microprocessor, and ASIC*. CRC Press, Boca Raton, FL.
- [14] E. Chung and J. Fowers et al. 2018. Serving DNNs in real time at datacenter scale with project brainwave. *IEEE Micro* 38, 2 (Mar. 2018), 8–20.
- [15] F. de Dinechin and A. Tisserand. 2001. Some improvements on multipartite table methods. In *Proceedings of the IEEE Symposium on Computer Arithmetic*. 128–135.
- [16] V. S. Dimitrov, G. A. Jullien, and W. C. Miller. 1999. Theory and applications of the double-base number system. *IEEE Trans. Comput.* 48, 10 (1999), 1098–1106. <https://doi.org/10.1109/12.805158>
- [17] Z. G. Feng and K. F. C. Yiu. 2010. A new formulation for the design of FIR filters with reduced hardware implementation complexity. In *Proceedings of the IEEE International Conference on Green Circuits Systems*. 242–246.
- [18] A. Garofalo, M. Rusci, F. Conti, D. Rossi, and L. Benini. 2019. PULP-NN: accelerating quantized neural networks on parallel ultra-low-power RISC-V processors. *Philos. Trans. Royal Soc. A* 378, 2164 (Dec. 2019). <https://doi.org/10.1098/rsta.2019.0155>
- [19] I. Hubara, M. Courbariaux, D. Soudry, R. El-Yaniv, and Y. Bengio. 2017. Quantized neural networks: Training neural networks with low-precision weights and activations. *ACM J. Mach. Learn. Res.* 18, 1 (Jan. 2017), 6869–6898.
- [20] K. Johansson, O. Gustafsson, and L. Wanhammar. 2008. Implementation of elementary functions for logarithmic number systems. *IET Comput. Digit. Tech.* 2, 4 (June 2008), 295–304.

- [21] N. G. Kingsbury and P. J. W. Rayner. 1971. Digital filtering using logarithmic arithmetic. *Electron. Lett.* 7, 2 (Jan. 1971), 56–58.
- [22] I. Kouretas, Ch. Basetas, and V. Paliouras. 2013. Low-power logarithmic number system addition/subtraction and their impact on digital filters. *IEEE Trans. Comput.* 62, 11 (2013), 2196–2209.
- [23] I. Kouretas and V. Paliouras. 2018. Logarithmic number system for deep learning. In *Proceedings of the International Conference on Modern Circuits and Systems Technologies*. 1–4.
- [24] V. B. Lawrence and A. C. Salazar. 1980. Finite precision design of linear-phase FIR filters. *Bell Syst. Tech. J.* 59, 9 (Nov. 1980), 1575–1598. <https://doi.org/10.1002/j.1538-7305.1980.tb03051.x>
- [25] E. H. Lee, D. Miyashita, E. Chai, B. Murmann, and S. S. Wong. 2017. LogNet: Energy-efficient neural networks using logarithmic computation. In *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing*. 5900–5904.
- [26] A. M. Mansour, A. M. El-Sawy, M. S. Aziz, and A. T. Sayed. 2015. A new hardware implementation of base 2 logarithm for FPGA. *Int. J. Signal Proc. Syst.* 3, 2 (2015), 171–181.
- [27] Alan Mishchenko. 2005. ABC: A System for Sequential Synthesis and Verification. Retrieved from <https://people.eecs.berkeley.edu/~alanmi/abc/>.
- [28] J. N. Mitchell. 1962. Computer multiplication and division using binary logarithms. *IRE Trans. Electron. Comp.* EC-11, 4 (1962), 512–517.
- [29] Sanjit K. Mitra. 2006. *Digital Signal Processing*. TATA McGraw-Hill, University of California, Santa Barbara, CA.
- [30] D. Miyashita, E. H. Lee, and B. Murmann. 2016. Convolutional neural networks using logarithmic data representation. Retrieved from <https://arXiv:cs.NE/1603.01025v2>.
- [31] B. Nam, H. Kim, and H. Yoo. 2008. Power and area-efficient unified computation of vector and elementary functions for handheld 3D graphics systems. *IEEE Trans. Comput.* 57, 4 (2008), 490–504.
- [32] V. Paliouras and T. Stouraitis. 2000. Logarithmic number system for low-power arithmetic. In *Proceedings of the International Workshop on Power Timing Modeling Optimization Simulation*, Vol. 1918. 285–294.
- [33] M. Rastegari, V. Ordonez, J. Redmon, and A. Farhadi. 2016. XNOR-Net: ImageNet Classification Using Binary Convolutional Neural Networks. Retrieved from <https://arXiv:cs.CV/1603.05279v4>.
- [34] M. Satyanarayanan. 2017. The emergence of edge computing. *Computer* 50, 1 (Jan. 2017), 30–39.
- [35] A. S. Sedra, K. C. Smith, T. C. Carusone, and V. Gaudet. 2019. *Microelectronic Circuits* (8th ed.). Oxford University Press.
- [36] W. Shi, J. Cao, Q. Zhang, Y. Li, and L. Xu. 2016. Edge computing: Vision and challenges. *IEEE Internet Things J.* 3 (2016), 637–646. Issue 5. <https://doi.org/10.1109/JIOT.2016.2579198>
- [37] G. Sicuranza. 1982. On the accuracy of 2-D digital filter realizations using logarithmic number systems. In *Proc. IEEE Int. Conf. Acoust. Speech Signal Process.*, Vol. 7. 48–51. <https://doi.org/10.1109/ICASSP.1982.1171386>
- [38] T. Stouraitis and V. Paliouras. 2002. Considering the alternatives in low-power design. *IEEE Circ. Syst. Mag.* 17, 4 (Aug. 2002), 22–29.
- [39] X. Sun, J. Choi, C. Y. Chen, N. Wang, S. Venkataramani, V. Srinivasan, X. Cui, W. Zhang, and K. Gopalakrishnan. 2019. Hybrid 8-bit floating point (HFP8) training and inference for deep neural networks. In *Proceedings of the Conference on Advances in Neural Information Processing Systems*. 4901–4910.
- [40] E.E. Swartzlander Jr. and A.G. Alexopoulos. 1975. The sign/logarithm number system. *IEEE Trans. Comput.* C-24, 12 (Dec. 1975), 1238–1242.
- [41] J. N. Coleman *et al.* 2008. The european logarithmic microprocessor. *IEEE Trans. Comput.* 57, 4 (Apr. 2008), 532–546.
- [42] Sebastian Vogel, Mengyu Liang, Andre Guntoro, Walter Stechele, and Gerd Ascheid. 2018. Efficient hardware acceleration of CNNs using logarithmic data representation with arbitrary log-base. In *Proceedings of the IEEE/ACM Conference on Computer-Aided Design (ICCAD’18)*. ACM, New York, NY, Article 9, 8 pages.
- [43] N. Wang, J. Choi, D. Brand, C.-Y. Chen, and K. Gopalakrishnan. 2018. Training deep neural networks with 8-bit floating point numbers. In *Proceedings of the Conference on Advances in Neural Information Processing Systems*. 7675–7684.
- [44] Y. Wang, J. Lin, and Z. Wang. 2018. An energy-efficient architecture for binary weight convolutional neural networks. *IEEE Trans. VLSI Syst.* 26, 2 (Feb. 2018), 280–293.
- [45] B.-D. Yang and L.-S. Kim. 2003. A low-power charge-recycling ROM architecture. *IEEE Trans. VLSI Syst.* 4, 4 (2003), 590–600.
- [46] M. Zhu, Y. Ha, C. Gu, and L. Gao. 2016. An optimized logarithmic converter with equal distribution of relative errors. *IEEE Trans. Circ. Syst. II: Express Briefs* 63, 9 (2016), 848–852.

Received September 2020; revised April 2021; accepted April 2021