Citation	Nimish Shah, Laura I. Galindez Olascoaga, Wannes Meert, Marian Verhelst, (2019), ProbLP: A framework for low-precision probabilistic inference Proceedings of the 56th Annual Design Automation Conference 2019, 190:1-190:6.
Archived version	Author manuscript: the content is identical to the content of the published paper, but without the final typesetting by the publisher
Published version	http://doi.acm.org/10.1145/3316781.3317885
Journal homepage	https://dl.acm.org/dl.cfm
Author contact	Nimish.shah@esat.kuleuven.be + 32 (0)16 321906

(article begins on next page)

PROBLP: A framework for low-precision probabilistic inference

Nimish Shah

MICAS, Electrical Engineering, KU Leuven
nshah@esat.kuleuven.be

Wannes Meert

DTAI, Computer Science, KU Leuven
wannes.meert@kuleuven.be

Laura I. Galindez Olascoaga

MICAS, Electrical Engineering, KU Leuven
laura.galindez@kuleuven.be

Marian Verhelst

MICAS, Electrical Engineering, KU Leuven
marian.verhelst@kuleuven.be

ABSTRACT

Bayesian reasoning is a powerful mechanism for probabilistic inference in smart edge-devices. During such inferences, a low-precision arithmetic representation can enable improved energy efficiency. However, its impact on inference accuracy is not yet understood. Furthermore, general-purpose hardware does not natively support low-precision representation. To address this, we propose PROBLP, a framework that automates the analysis and design of low-precision probabilistic inference hardware. It automatically chooses an appropriate energy-efficient representation based on worst-case error-bounds and hardware energy-models. It generates custom hardware for the resulting inference network exploiting parallelism, pipelining and low-precision operation. The framework is validated on several embedded-sensing benchmarks.

KEYWORDS

Embedded machine learning, Energy efficiency, Probabilistic inference, Arithmetic circuits, Bayesian networks, Sum product networks, Low-precision, Error bounds

1 INTRODUCTION

The use of probabilistic inference is popular for robust classification, diagnosis and decision-making problems, because of its ability to assign a confidence-level to every result, in terms of probability. Probabilistic Graphical Model (PGM) [16], an established tool for probabilistic inference, is widely used for such problems. PGMs have several interesting properties that make them suitable for embedded applications. Specifically, PGMs: 1) are capable of dealing with missing data; 2) allow to incorporate information from different domains, as well as expert knowledge; 3) can be trained with less data; and 4) can explicitly model uncertainty and causal relationships in the system. In addition, PGMs' performance is competitive with respect to other state-of-the-art Machine Learning implementations on embedded sensing applications [9, 19, 10, 13].

Inference in PGMs is prominently performed using a versatile representation known as an Arithmetic Circuit (AC) (or sum-product network) [6]. An AC refers to a model of computation and is often represented as a graph of additions and multiplications. ACs allow for an integration of both statistical and symbolic methods in artificial intelligence, a promising combination that is pursued in state-of-the-art machine learning methods [17, 14, 13]. They are also central to performing inference in the field of probabilistic (logic) programming [8, 14]. Furthermore, recent approaches learn

ACs directly from data, with state-of-the-art performance in certain applications [13]. In this work, we focus on ACs representing Bayesian networks (BN), a type of PGM.

Inference in ACs is generally restricted to obtaining exact solutions on general purpose computing devices. A significant improvement in energy efficiency would be possible by tolerating some error through approximating the probability computed by ACs. Take for instance a smartphone-based activity identification for elderly, wherein a probability is evaluated for different activities (e.g., a user walking up the stairs). The application chooses to identify an activity only if its probability is higher than a certain threshold, say 0.60. Here, allowing an output error of 0.01 would only affect the decisions within the probability range of 0.59 and 0.61, while enabling improved energy-efficiency.

A promising hardware optimization that can exploit the available error-tolerance is to realize the additions and multiplications in reduced-precision representation. Yet, the state-of-the-art is lacking analysis of the impact of such precision reduction on the output probability error. Previous works [5, 4, 18] have studied the impact of low-precision in leaf nodes of an AC, but do not account for noisy or low precision computations in its internal nodes. However, the error in the imprecise internal nodes can get accumulated and be the dominant source of imprecision in the inference output.

In this paper, we propose PROBLP¹, a holistic framework to automate the design of low-precision energy-efficient hardware for probabilistic inference in Arithmetic Circuits. Our contributions are as follows:

- We derive bounds on the error in probabilistic queries due to low-precision representation, taking into account the error introduced in all the nodes in an AC.
- We develop energy models to help choose the most energy-efficient representation
- We develop a tool to automatically generate low-precision inference hardware, and validates its performance on several embedded sensing benchmarks.

This paper is organized as follows. Section 2 gives an introduction to Arithmetic circuits compiled from Bayesian Networks and an overview of related works. In Section 3, we derive analytical error bounds for ACs. We introduce the PROBLP framework and elaborate on how it selects the optimal precision and selects between fixed- or floating-point representation. Section 4 demonstrates the validity of the framework on a suite of embedded sensing benchmarks and Section 5 concludes this work.

¹Code available at <https://github.com/nimish15shah/ProBLP>

Figure 1: (a) Example of a Bayesian Network and (b) Arithmetic compiled from the BN of Figure 1a.

2 BACKGROUND AND PREVIOUS WORK

In this paper, we denote a random variable with uppercase letter X and its instantiation with a lower case letter x . A set of multiple random variables are denoted with bold upper case letters $\mathbf{X}$ and its joint assignment with bold lower case letters $\mathbf{x}$.

Bayesian networks (BN) are directed acyclic graphs that compactly encode a joint probability distribution over a set of random variables $\{X_1, \dots, X_n\}$ [16]:

$$Pr(X_1, \dots, X_n) = \prod_{i=1}^n Pr(X_i | \Pi_{X_i}), \quad (1)$$

where Π_{X_i} denotes the parents of X_i and $Pr(X_i | \Pi_{X_i})$ are the conditional dependencies between variables and their parents, which can be represented as Conditional Probability Tables (CPTs). In the graphical component of BNs, the variables are represented as nodes and their probabilistic or causal relationships are indicated by the direction of the edges among them, as depicted in Figure 1a. The joint probability distribution in (1) allows answering a number of probabilistic queries such as the marginal probability, the conditional probability or the Most Probable Explanation (MPE)[16].

Probabilistic inference on a BN can be made efficient by compiling it to an Arithmetic circuit, which consists only of multiplications and addition. Figure 1b shows an example of an AC generated by compiling the BN in Figure 1a. The inputs to this AC are the BN's parameters, represented by $\theta_{x|u}$, where x and u are the instantiation of random variable X and its parents. The second type of inputs to the AC are indicators λ_x , which are binary variables that indicate the evidence of the observed nodes. The probability of an evidence (e.g., $\mathbf{e} = \{A = a_1, C = c_3\}$) can be computed with an upward pass on the AC by setting the indicators that contradict the evidence to 0 ($\lambda_{a_2} = \lambda_{c_1} = \lambda_{c_2} = 0$), and others to 1 ($\lambda_{a_1} = \lambda_{b_1} = \lambda_{b_2} = \lambda_{c_3} = 1$).

Previous works have studied the impact of finite-precision CPT parameters on marginal and conditional probability [5, 4, 18]. This helps to reduce memory footprint due to a smaller inference model. However, they did not study the effect of low-precision arithmetic operations. The authors of [20, 12] studied the effect of fixed-point

arithmetic on marginal probability for a few BNs, but not of conditional probability, and did not provide error bounds. Moreover, the impact under floating-point arithmetic operation is also unclear.

This work provides analytical bounds on the absolute and the relative error in marginal and conditional probabilities for fixed- and floating-pt operations in the entire AC. A holistic framework PROBLP is introduced, which also takes energy-consumption into account to choose the optimal representation among fixed-point and floating-point. Subsequently, it automatically generates custom hardware for the AC evaluation.

3 METHODOLOGY

Different components of the PROBLP framework are shown in figure 2. PROBLP takes in an AC together with some user requirements, based on which, it calculates the least number of fixed and floating point bits needed to meet these requirements. To do so, PROBLP evaluates error-bounds for the AC, based on their error models. To choose between these two representations, it subsequently estimates the energy of the complete AC based on energy models. Finally, it generates a fully-parallel pipelined hardware in the selected low-precision representation.

The three inputs of PROBLP are as follows:

Arithmetic circuit: The Arithmetic circuit to be implemented using low-precision hardware. In this paper, we use ACs compiled from Bayesian networks, but they can as well be compiled from probabilistic (logic) programs or can be trained directly from data.

Type of query: The type of probabilistic query to be performed using the AC, to be chosen from *marginal probability*, *conditional probability* or the probability of *most probable explanation* (MPE).

Error tolerance: The amount of error on the output that can be tolerated in the probabilistic queries by the application, for *all* possible combination of inputs, in terms of *absolute* or *relative* error. An absolute error is given as $\epsilon = \tilde{Pr} - Pr$ and a relative error is given as $\epsilon = \frac{\tilde{Pr} - Pr}{Pr}$, where Pr is the output probability of interest.

3.1 Error analysis

The aim of error analysis is to estimate the minimum number of bits required to achieve the user-specified error tolerance. For this, it has to take into account the impact of reducing the number of bits on the error in the AC output probability. There are two sources of error in an AC: an error in the leaf nodes when CPT values are quantized to finite precision, and an error injected in the intermediate nodes of AC due to the finite precision arithmetic operations. Unlike previous research works, we formally treat the error in the intermediate nodes as well to derive the error-bounds.

We consider two representations: fixed-point and floating-point. Operators of both types round bits during computation. For example, a multiplication of 2 inputs of n bits produces an exact result with $2n$ bits, which is subsequently rounded to fit an n bit output. The error introduced can be modeled as an additive noise source.

The error models used for the leaf nodes and the intermediate nodes are described next. Some of the models are inspired by [11], but authors of that work did not perform the error analysis for ACs, and some of the models would render unbounded errors if AC-specific constraints are not exploited. The arithmetic operators are assumed to round the extra bits to the nearest value.

Figure 2: Internal working of PROBLP

3.1.1 Fixed-pt error estimation. Let I and F be the number of integer and fraction bits. All the numbers are assumed to be in the range of the fixed-pt format, implying an absence of overflow during computation, which can be ensured by using an appropriate number of integer bits I , discussed in detail in section 3.1.4.

Fixed-pt leaf node: Let a be the real value of a leaf node in an AC, and $\tilde{a}$ be its fixed-pt representation. The error in fixed-pt conversion can be bounded as,

$$|\Delta a| = |\tilde{a} - a| \leq 2^{-(F+1)} \quad (2)$$

Fixed-pt adder node: If $\tilde{a}$ and $\tilde{b}$ are the fixed-pt representation of adder inputs a and b , the error in the output $\tilde{f}$ is given as

$$\begin{aligned} \Delta f = \tilde{f} - f &= (\tilde{a} + \tilde{b}) - (a + b) = (a + \Delta a + b + \Delta b) - (a + b) \\ &= \Delta a + \Delta b \end{aligned} \quad (3)$$

Note that the fixed-pt adder does not add any error of its own, as it does not round bits, and hence simply accumulates the error of the inputs. Note again that the adder output cannot overflow, as all the numbers are ensured to be in range.

Fixed-pt multiplier node: With $\tilde{a}$ and $\tilde{b}$ as the fixed-pt representation of multiplier inputs a and b , the error in fixed-pt multiplier output $\tilde{f}$ can be bounded as,

$$\Delta f = \tilde{f} - f \leq (\tilde{a}\tilde{b} + 2^{-(F+1)}) - ab \quad (4)$$

$$\begin{aligned} &= (a + \Delta a)(b + \Delta b) - ab + 2^{-(F+1)} \\ &\leq a_{\max}\Delta b + b_{\max}\Delta a + \Delta a\Delta b + 2^{-(F+1)} \end{aligned} \quad (5)$$

In (4), the error term $2^{-(F+1)}$ models the error introduced when the LSB bits of the intermediate multiplication result are rounded to fit back into F fractional bits. Equation (5) produces an unbounded error unless $a_{\max}$ and $b_{\max}$ can be bounded.

The $a_{\max}$ and $b_{\max}$ can be efficiently bounded by taking into account the AC-specific properties. An AC consists of adders and multipliers and only operates on non-negative numbers. As a result, each internal node in the AC is a monotonously increasing function of its inputs. Hence, all the nodes are at the maximum value when all the inputs are at their maximum. As such, since CPT parameters stay constant across AC evaluations, this is achieved when all the indicator variables λ_x are set to 1. This allows to assess the $a_{\max}$ and

$b_{\max}$ of every operator in the AC with just a single AC evaluation. Thereby, allowing PROBLP to bound the error of fixed-pt multipliers.

3.1.2 Floating-pt error estimation. Let E and M be the exponent and mantissa bits. We only consider normalized floating-pt here. All the numbers are assumed to be within the range of the given format, ensured by a method explained in detail in section 3.1.4.

Float-pt leaf node: Let a be the real value of a leaf in AC and $\tilde{a}$ be its floating-pt representation. The absolute error introduced due to the floating-pt conversion can be bounded as described in [11],

$$\left| \frac{\Delta a}{a} \right| \leq 2^{-(M+1)} \quad (6)$$

which can be expressed alternatively as,

$$\tilde{a} = a + \Delta a = a \left(1 + \frac{\Delta a}{a} \right) = a(1 \pm \epsilon) \quad (7)$$

where:

$$0 \leq \epsilon \leq 2^{-(M+1)}$$

Float-pt adder node: Let $\tilde{a}$ and $\tilde{b}$ be the float-pt versions of adder inputs a and b , f be the ideal output, and $\tilde{f}$ be the output of a floating-pt adder. $\tilde{a}$ and $\tilde{b}$ can be represented as,

$$\tilde{a} = a(1 \pm \epsilon)^m \quad \text{and} \quad \tilde{b} = b(1 \pm \epsilon)^n \quad (8)$$

Here, m and n depends on the amount of error accumulated in a and b , respectively. The bound on $\tilde{f}$ can be given as follows,

$$\tilde{f} = (\tilde{a} + \tilde{b})(1 \pm \epsilon) = (a(1 \pm \epsilon)^m + b(1 \pm \epsilon)^n)(1 \pm \epsilon) \quad (9)$$

$$\begin{aligned} &= (a + b) \left(\frac{a}{a+b} (1 \pm \epsilon)^m + \frac{b}{a+b} (1 \pm \epsilon)^n \right) (1 \pm \epsilon) \\ &= f \left(\frac{a}{a+b} (1 \pm \epsilon)^{m+1} + \frac{b}{a+b} (1 \pm \epsilon)^{n+1} \right) \\ &\leq \begin{cases} f(1 \pm \epsilon)^{m+1} & m > n \\ f(1 \pm \epsilon)^{n+1} & \text{otherwise} \end{cases} \end{aligned} \quad (10)$$

The error term in (9) is due to the rounding of LSB bits of the mantissa of the smaller input before addition.

Float-pt multiplier node: Just as in case of the adder, the inputs $\tilde{a}$ and $\tilde{b}$ can be bounded as in (8). With that, the output of a floating-pt multiplier can be given as,

$$\tilde{f} = \tilde{a}\tilde{b}(1 \pm \epsilon) = ab(1 \pm \epsilon)^m(1 \pm \epsilon)^n(1 \pm \epsilon) \quad (11)$$

$$= f(1 \pm \epsilon)^{m+n+1} \quad (12)$$

The error term in (11) is due to the rounding of the LSB bits of the mantissa to fit the result in M mantissa bits.

3.1.3 Error-bound at the AC output. Equations (2), (3), (5), and (6), (10) and (12) corresponds to the *Error models* shown in figure 2. As these models generate the output in the same format as the inputs, they provide a way to recursively propagate the error from the leaves of an AC all the way up to its output node, by accumulating the error introduced in every adder and multiplier. Figure 3 shows an example of error propagation using the fixed-pt error models. This is performed as a part of the *fixed-pt error analysis* and *float-pt error analysis* blocks of PROBLP shown in figure 2.

Figure 3: Error propagation with fixed-pt error models.

The error propagation in fixed-pt arithmetic produces a bound of the form $\Delta f \leq c$, where Δf is the absolute error in the output node, and c is a constant that depends on the size and structure of the AC, its parameters, and the number of fixed-pt bits. The constant c can be estimated recursively with our error models for any given AC. Similarly, the error propagation in floating-pt arithmetic produces a bound of the form $\tilde{f} \leq f(1 \pm \epsilon)^c$, where $\tilde{f}$ is the output of an AC with floating-pt operators, f is the ideal output, ϵ is a constant related to the number of floating-pt bits, and c is a constant related to the size and structure of the AC. Again, the constant c can be estimated recursively using the models we proposed, for any AC. Alternatively, the floating-pt bound can be expressed as $\frac{\tilde{f}-f}{f} \leq \delta$ for some constant δ , i.e., a bounded relative error at the output.

3.1.4 Number of integer or exponent bits. For the error-models proposed in section 3.1.1 and 3.1.2 to be valid, the numbers encountered during the computation must stay within range of the representation. This can be ensured by using an appropriate number of integer bits I and exponent bits E for fixed- and floating-pt respectively. Otherwise, error in some of the probability evaluations would exceed the predicted bounds. It is hence important to automatically derive the required range of numbers for any given AC.

Max-value analysis: The largest number to be encountered in an AC can be derived by setting all the indicator variables λ_x to 1, as explained in section 3.1.1. Analyzing the internal AC data values of this query, allows deriving the required I , resp. E to avoid overflow.

Min-value analysis: The floating-pt models are invalid in case of underflow as well. Hence, it is necessary to also estimate the smallest positive non-zero value for an AC. It can be proven that all the nodes in an AC are at the respective minimum non-zero values when all the indicator variables λ_x are set to 1 and the adders are replaced with minimum operators $\min(i, j)$. The resulting efficient AC evaluation allows PROBLP to analyze a lower bound on AC values, and find the appropriate E required to prevent underflow. The fixed-pt models remain valid even if the number of fraction bits is not enough to represent small values in the AC, so no special precautions are needed here regarding underflow.

In this way, PROBLP performs the Max-value and Min-value analysis to select E , resp. I , that satisfies both the requirements.

3.2 Bounds for probabilistic queries

As shown in figure 2, PROBLP aims to estimate the optimal fixed-pt and float-pt bit width for a given type of probabilistic query and error tolerance. However, the bounds derived so far, apply only to a single AC evaluation. Some types of probabilistic queries require a combination of multiple AC evaluations. In this section,

we derive bounds for two type of probabilistic queries: 1) Marginal probability and MPE, and 2) Conditional probability.

3.2.1 Marginal probability and MPE. Marginal probabilities $Pr(\mathbf{q}, \mathbf{e})$ and most probable explanation (MPE) need only one AC evaluation. Hence, the bounds derived in section 3.1.3 apply for these queries.

3.2.2 Conditional probability. Conditional probability $Pr(\mathbf{q} | \mathbf{e})$ is evaluated by performing two AC evaluations, one for $Pr(\mathbf{q}, \mathbf{e})$ and one for $Pr(\mathbf{e})$, followed by taking the ratio of the two results².

Fixed-pt bounds: In the case of fixed-pt arithmetic, the absolute error in each of the AC queries remains bounded. The impact on the conditional probability can hence be given as,

$$\tilde{Pr}_{fx}(\mathbf{q} | \mathbf{e}) = \frac{\tilde{Pr}_{fx}(\mathbf{q}, \mathbf{e})}{\tilde{Pr}_{fx}(\mathbf{e})} = \frac{Pr(\mathbf{q}, \mathbf{e}) + \Delta_1}{Pr(\mathbf{e}) + \Delta_2} \quad (13)$$

Here, maximum error is achieved when $\Delta_2 = 0$ and $\Delta_1 = \Delta_{1max}$. In such a case, following equations show the impact on absolute and relative error in the conditional probability,

$$\begin{aligned} \tilde{Pr}_{fx}(\mathbf{q} | \mathbf{e}) &= Pr(\mathbf{q} | \mathbf{e}) + \frac{\Delta_{1max}}{Pr(\mathbf{e})} \\ \Rightarrow \Delta Pr_{fx}(\mathbf{q} | \mathbf{e}) &= \frac{\Delta_{1max}}{Pr(\mathbf{e})} \leq \frac{\Delta_{1max}}{\min Pr(\mathbf{e})} \end{aligned} \quad (14)$$

$$\Rightarrow \frac{\Delta Pr_{fx}(\mathbf{q} | \mathbf{e})}{Pr(\mathbf{q} | \mathbf{e})} = \frac{\Delta_{1max}}{Pr(\mathbf{e}) Pr(\mathbf{q} | \mathbf{e})} \quad (15)$$

Equations (14) and (15) show the absolute and relative error in the conditional query $Pr(\mathbf{q} | \mathbf{e})$. The error in the numerator is scaled by $Pr(\mathbf{e})$ and $Pr(\mathbf{e}) Pr(\mathbf{q} | \mathbf{e})$, and these probabilities can become very small. Hence, large number of fixed-pt bits are generally required to achieve a reasonable error-bound, especially for the relative-error bound of (15). The absolute-error bound in (14) can be quantified by estimating the minimum possible value for $Pr(\mathbf{e})$ as described in section 3.1.4, wherein adders are replaced with min operators.

As the denominator of (15) can become very small, it is not a good idea to use fixed-pt when requiring a relative error bound in conditional probabilities. Moreover, quantifying a bound for (15) is also not straightforward. Hence, PROBLP will always choose float-pt for relative error in conditional probability.

Float-pt bounds: The impact of using float-pt arithmetic on conditional probability can be given as follows.

$$\begin{aligned} \tilde{Pr}_{fl}(\mathbf{q} | \mathbf{e}) &= \frac{\tilde{Pr}_{fl}(\mathbf{q}, \mathbf{e})}{\tilde{Pr}_{fl}(\mathbf{e})} = \frac{Pr(\mathbf{q}, \mathbf{e})(1 \pm \epsilon)^m}{Pr(\mathbf{e})(1 \pm \epsilon)^n} \\ &= Pr(\mathbf{q} | \mathbf{e})(1 \pm \epsilon)^{m-n} \end{aligned} \quad (16)$$

In (16), m and n are upper bounded by a constant, say c , but not lower bounded. In the worst case, one of them can become 0, while the other is c . Even in this worst case, the floating-pt version of the conditional probability still remains bounded as follows.

$$Pr(\mathbf{q} | \mathbf{e})(1 \pm \epsilon)^{-c} \leq \tilde{Pr}_{fl}(\mathbf{q} | \mathbf{e}) \leq Pr(\mathbf{q} | \mathbf{e})(1 \pm \epsilon)^c \quad (17)$$

This ensures a bound on the relative error $\frac{\Delta Pr_{fl}(\mathbf{q} | \mathbf{e})}{Pr(\mathbf{q} | \mathbf{e})}$.

² $Pr(\mathbf{q} | \mathbf{e})$ can also be estimated by an upward and a downward pass in an AC followed with a division. We do not consider it explicitly, but similar error bounds are expected.

3.3 Selecting optimal representation

Section 3.1 and 3.2 establishes a method to evaluate error bounds for a given AC in terms of number of bits. Next, PROBLP finds the least number of fixed-pt and float-pt bits needed for given requirements. For this, it evaluates the bounds starting with 2 fraction bits and 2 mantissa bits, and increments them until the error-requirement is satisfied. Then, it estimates the least number of integer and exponent bits required by the min and max analysis explained in section 3.1.4. In this way, PROBLP comes up with the optimal fixed-pt and float-pt representation shown in figure 2.

Subsequently, the framework has to select among fixed-pt and float-pt. PROBLP selects the one with the lowest energy-consumption, estimated using operator-level energy models. Energy models for the adders and multipliers are developed by synthesizing them with varying fraction/mantissa bits and integer/exponent bits in TSMC 65nm technology and extracting post-synthesis energy consumption. The models were fitted using least-squares method to the simulation results, and are summarized in Table 1.

Table 1: Energy models for arithmetic operators at 1V. N is the #fixed-pt bits and M is the #mantissa bits

Operator	Energy (fJ)
Fixed-pt add	7.8 N
Fixed-pt mult	$1.9 N^2 \log N$
Float-pt add	$44.74 (M+1)$
Float-pt mul	$2.9 (M+1)^2 \log (M+1)$

3.4 Automatic hardware generation

PROBLP suggests the most-appropriate low-precision representation for the AC, but this may not translate to energy savings unless hardware has custom arithmetic operators. To address this, PROBLP has an integrated hardware generator that generates custom parallel hardware that is fully-pipelined and consists of arithmetic operators of the exact precision that is required to meet the user requirements. There are two major stages in the hardware generation process. In the first stage, all AC operators with more than two inputs are decomposed into a tree of 2-input operators. An example of such decomposition is shown in figure 4, wherein the F operator is decomposed into a tree of F1, F2 and F3. In the second stage, the generator inserts pipeline registers after every operator. In some cases, it may have to insert multiple registers due to a mismatch in path timings, as shown in the path between A and G in figure 4. The final output of PROBLP is a verilog code of the custom hardware.

4 EXPERIMENTAL RESULTS

We validate the functionality of PROBLP for the Arithmetic circuits targeting embedded sensing applications, by performing two types of experiments on four datasets. Three of these datasets (HAR [1], UNIMIB[15], UIWADS [3] in table 2) correspond to activity and user identification applications in smartphones and therefore rely on the accurate estimation of a conditional probability of the form $Pr(Activity|sensors)$ to make threshold-based decisions. The fourth dataset (Alarm in table 2 [2]) is of a patient monitoring application and is often used as a standard Bayesian network benchmark.

Figure 4: Automatic conversion of an AC to a pipelined hardware.

The ACs used in this section are compiled using the ACE tool [7], with -cd06 and -forceC2d option enabled. For the experiments on HAR, UNIMIB, and UIWADS, we trained Naive Bayes classifier on 60% of the data and used the rest for testing. The testing dataset for Alarm is generated by sampling 1000 instances from the trained network. In all the experiments, the leaf nodes of the BN were used as evidence nodes e and one of the root nodes in the BN (the class node in the case of the classifiers) as a query node q .

4.1 Validation of bounds

This experiment confirms the validity of the derived error bounds for the AC compiled from the Alarm network. The experimental setting is as follows:

Fixed-pt: The number of integer bits is set to 1 based on the max-analysis, and fraction bits is varied from 8 to 40.

Float-pt: The number of exponent bits is set to 8 based on the max-min analysis, and mantissa bits is varied from 8 to 40.

Figure 5 shows the max and mean error on the test-set, which confirm the validity of the bounds.

4.2 Overall performance

In this experiment, the complete PROBLP framework is deployed to choose an appropriate arithmetic representation and generate hardware for different ACs and for given user requirements. The results of the experiment are summarized in Table 2. Experiments are performed for all combinations of queries and types of error tolerances for the HAR AC, and two combinations for the rest of the ACs. The table shows the optimal fixed-pt and float-pt representation that meet the target error-tolerance. Among these, PROBLP selects the one with less predicted energy, highlighted in bold. The resulting maximum error observed on the test-sets remain within the required error-tolerance. The post-synthesis energy consumption matches well to the energy predicted by the framework. The energy consumption of the hardware with a 32b float ($E=8$, $M=23$, 1 sign bit) is also shown for comparison. Note here that the choice of 0.01 error tolerance is arbitrary and higher energy-efficiency can be achieved for relaxed error tolerances.

5 CONCLUSION

Probabilistic inference with Arithmetic circuits can be made energy-efficient by tolerating a small amount of error in output probabilities and by designing custom hardware to exploit this error tolerance.

Table 2: Optimal fixed-pt and float-pt representations that meets error tolerance with selected repr. in bold. Measured error and energy for the selected representation. I,F,E and M stands for number of integer, fraction, exponent and mantissa bits

AC	Type of query	Error tolerance	Opt. Fx-pt repr. I, F (pred. energy in nJ/AC_eval)	Opt. Fl-pt repr. E, M (pred. energy in nJ/AC_eval)	Max error observed on test-set	Post-synth. energy (nJ/AC_eval)	Energy of 32b Fl-pt E=8, F=23
HAR	Marg. prob.	abs. err 0.01	1, 15 (4.3)	9, 14 (6.7)	5.9×10^{-4}	5.3	10.8
	Marg. prob.	rel. err 0.01	1, >64 (-)	9, 14 (6.7)	1.0×10^{-3}	7.2	
	Cond. prob.	abs. err 0.01	1, >64 (-)	9, 14 (6.7)	2.6×10^{-4}	7.2	
	Cond. prob.	rel. err 0.01	-	9, 14 (6.7)	1.0×10^{-3}	7.2	
UNIMIB	Marg. prob.	abs. err 0.01	1, 13 (0.4)	7, 12 (0.6)	4.9×10^{-4}	0.34	0.89
	Cond. prob.	rel. err 0.01	-	7, 12 (0.6)	1.1×10^{-3}	0.44	
UIWADS	Marg. prob.	abs. err 0.01	1, 11 (0.06)	6, 10 (0.09)	1.3×10^{-3}	0.06	0.18
	Marg. prob.	rel. err 0.01	1, 47 (1.3)	6, 10 (0.09)	1.2×10^{-3}	0.08	
Alarm	Marg. prob.	abs. err 0.01	1, 14 (2.2)	8, 13 (3.2)	2.2×10^{-4}	2.43	5.37
	Cond. prob.	rel. err 0.01	-	8, 13 (3.2)	2.8×10^{-4}	3.18

Figure 5: Analytical error bounds and the observed error over a test-set, for varying (a) fraction and (b) mantissa bits, validating the correctness of bounds

This paper, therefore, proposes PROBLP, a holistic framework to automate the design of low-precision custom hardware for ACs. The framework estimates worst-case error bounds for ACs, taking into account the error incurred in reduced precision fixed- and floating-point operators. It estimates the impact of these errors on different types of probabilistic queries and finds the least number of fixed-pt and float-pt bits required to meet the error-tolerance. Subsequently, it chooses among the fixed-pt and float-pt representation based on the energy models developed for this purpose. Next, PROBLP automatically converts an AC to pipelined logic with custom arithmetic operators. The analytically derived error bounds are validated for varying fixed- and float-pt bits. Finally, the PROBLP framework is used for several embedded sensing benchmarks, confirming that the error-requirements are met and the energy consumption of automatically generated hardware matches the prediction.

ACKNOWLEDGMENT

This work has been supported by the EU ERC project Re-SENSE under grant agreement ERC-2016-STG-715037

REFERENCES

- [1] Davide Anguita, Alessandro Ghio, Luca Oneto, Xavier Parra, and Jorge Luis Reyes-Ortiz. 2013. A public domain dataset for human activity recognition using smartphones. In *ESANN*.
- [2] Ingo A Beinlich, Henri Jacques Suermondt, R Martin Chavez, and Gregory F Cooper. 1989. The alarm monitoring system: a case study with two probabilistic inference techniques for belief networks. In *AIME 89*. Springer, 247–256.
- [3] Pierluigi Casale, Oriol Pujol, and Petia Radeva. 2012. Personalization and user verification in wearable systems using biometric walking patterns. *Personal and Ubiquitous Computing*, 16, 5, 563–580.
- [4] Hei Chan and Adnan Darwiche. 2004. Sensitivity analysis in bayesian networks: from single to multiple parameters. In *Proceedings of the 20th conference on Uncertainty in artificial intelligence*. AUAI Press, 67–75.
- [5] Hei Chan and Adnan Darwiche. 2002. When do numbers really matter? *Journal of artificial intelligence research*, 17, 265–287.
- [6] Mark Chavira and Adnan Darwiche. 2008. On probabilistic inference by weighted model counting. *Artificial Intelligence*, 172, 6, 772–799. ISSN: 0004-3702. DOI: <https://doi.org/10.1016/j.artint.2007.11.002>. <http://www.sciencedirect.com/science/article/pii/S0004370207001889>.
- [7] Adnan Darwiche and Mark Chavira. [n. d.] Ace, an arithmetic circuit compiler. 2007. URL: <http://reasoning.cs.ucla.edu/ace>.
- [8] Daan Fierens, Guy Van den Broeck, Ingo Thon, Bernd Gutmann, and Luc De Raedt. 2015. Inference and learning in probabilistic logic programs using weighted CNFs. *Theory and Practice of Logic Programming*, 15.
- [9] Laura Galindez, Komail Badami, Jonas Vlasselaer, Wannes Meert, and Marian Verhelst. 2018. Dynamic sensor-frontend tuning for resource efficient embedded classification. *IEEE Journal on Emerging and Selected Topics in Circuits and Systems*.
- [10] Dileep George et al. 2017. A generative vision model that trains with high data efficiency and breaks text-based captchas. *Science*, 358, 6368, eaag2612.
- [11] Nicholas J Higham. 2002. *Accuracy and stability of numerical algorithms*. Vol. 80. Siam.
- [12] Osama U Khan and David D Wentzloff. 2016. Hardware accelerator for probabilistic inference in 65-nm cmos. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, 24, 3, 837–845.
- [13] Yitao Liang and Guy Van den Broeck. 2019. Learning logistic circuits. *Proceedings of the 33rd Conference on Artificial Intelligence (AAAI)*.
- [14] Robin Manhaeve, Sebastijan Dumančić, Angelika Kimmig, Thomas Demeester, and Luc De Raedt. 2018. Deepproblog: neural probabilistic logic programming. *NIPS*.
- [15] Daniela Micucci, Marco Mobilio, and Paolo Napoletano. 2017. Unimib shar: a dataset for human activity recognition using acceleration data from smartphones. *Applied Sciences*, 7, 10, 1101.
- [16] Judea Pearl. 1988. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann.
- [17] Clive Thompson. 2018. How to teach artificial intelligence some common sense. *Wired Magazine*, (Dec. 2018).
- [18] Sebastian Tschiatschek and Franz Pernkopf. 2015. On bayesian network classifiers with reduced precision parameters. *IEEE transactions on pattern analysis and machine intelligence*, 37, 4, 774–785.
- [19] Jonas Vlasselaer, Wannes Meert, and Verhelst Marian. 2018. Towards resource-efficient classifiers for always-on monitoring. In *Proceedings of the European Conference on Machine Learning 2018*.
- [20] Sara Zermani, Catherine Dezan, Hanen Chenini, Jean-Philippe Diguët, and Reinhardt Euler. 2015. Fpga implementation of bayesian network inference for an embedded diagnosis. In *Prognostics and Health Management (PHM), 2015 IEEE Conference on*. IEEE, 1–10.