

Acceleration of probabilistic reasoning through custom processor architecture

Nimish Shah*, Laura I. Galindez Olascoaga*, Wannes Meert[†] and Marian Verhelst*

*MICAS, Department of Electrical Engineering, KU Leuven, Belgium

[†]DTAI, Department of Computer Science, KU Leuven, Belgium

Email: {nimish.shah, laura.galindez, marian.verhelst}@esat.kuleuven.be, wannes.meert@cs.kuleuven.be

Abstract—Probabilistic reasoning is an essential tool for robust decision-making systems because of its ability to explicitly handle real-world uncertainty, constraints and causal relations. Consequently, researchers are developing hybrid models by combining Deep Learning with probabilistic reasoning for safety-critical applications like self-driving vehicles, autonomous drones, etc. However, probabilistic reasoning kernels do not execute efficiently on CPUs or GPUs. This paper, therefore, proposes a custom programmable processor to accelerate sum-product networks, an important probabilistic reasoning execution kernel. The processor has an optimized datapath architecture and memory hierarchy optimized for sum-product networks execution. Experimental results show that the processor, while requiring fewer computational and memory units, achieves a 12x throughput benefit over the Nvidia Jetson TX2 embedded GPU platform.

Index Terms—Sum-product networks, Arithmetic circuits, Custom processor, Probabilistic reasoning, GPU, acceleration

I. INTRODUCTION

Deep learning (DL) has shown remarkable success in various computer vision and natural language processing tasks. However, it suffers from a serious drawback of not explicitly handling real-world uncertainty, constraints and causal relations, which are crucial for deploying robust decision-making systems in practice. To address this, researchers are developing hybrid models by combining DL with probabilistic reasoning techniques, eg. random sum-product networks [8], probabilistic logic neural networks [10], etc.

To ease the development of such hybrid models, several software frameworks like Pyro, DeepProbLog, etc. have been developed. However, deploying these hybrid models in real applications raises demands of adequate throughput, latency, energy consumption, etc. While several efficient hardware platforms have been developed for DL, probabilistic reasoning kernels are still implemented on general-purpose CPUs which is energy-inefficient. To address this, we propose a customized processor architecture that efficiently accelerates probabilistic reasoning workloads. Specifically, in this work:

- GPU execution bottlenecks are identified by developing a highly-optimized CUDA implementation.
- A programmable processor architecture is designed to alleviate these bottlenecks.

This work has been supported by the EU ERC project Re-SENSE under grant agreement ERC-2016-STG-715037, and we acknowledge support from Intel.

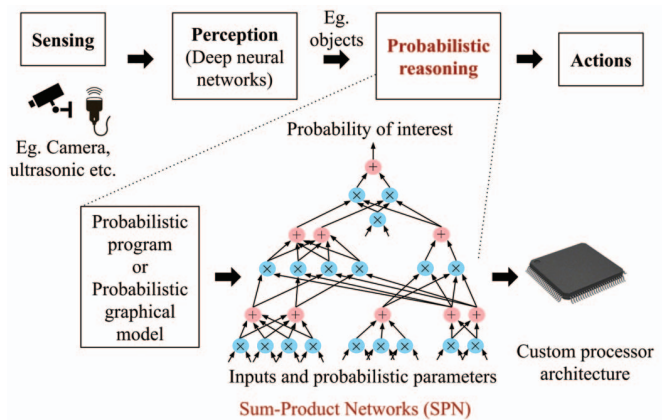


Fig. 1. An example of a hybrid system that uses Deep Learning for perception and probabilistic models for reasoning. Probabilistic models are usually converted to sum-product networks (SPN) for efficient inference.

- Experiments are performed to validate the performance improvement over a suite of benchmarks.

The paper is organized as follows: sec. II describes the compute kernel, while sec. III discusses its performance on a CPU and GPU. The proposed processor architecture is described in sec. IV. Sections V and VI evaluate the performance and discuss related work. Finally, sec. VII concludes the work.

II. PROBABILISTIC REASONING COMPUTE KERNEL

As shown in fig. 1, probabilistic reasoning is usually performed with probabilistic program or a probabilistic graphical model. These models can be converted to a tractable representation for efficient inference, called Sum-Product Network (SPN, also known as Arithmetic circuit) [1], [9]. Furthermore, SPNs can also be learned directly from data. SPNs are rooted directed acyclic graphs wherein the internal nodes are either sums or products, and the leaf nodes are probabilistic parameters or data inputs.

On a general-purpose processor, SPN can be implemented as a list of operations (alg. 1), or as a for loop over a vector (alg. 2). This loop formulation already shows that, although SPNs offer parallelism, they inherently consist of scalar operations that consume prior intermediate results via irregular indirect memory accesses. This leads to inefficient parallel execution as will be demonstrated in the next section.

Algorithm 1 SPN as a list of operations

Inputs: **IN:** SPN leaf nodes as a vector

- 1: $r0 = IN[0] \times IN[1]$
 - 2: $r1 = IN[2] \times IN[3]$
 - 3: $r2 = r0 + r1$
 - 4: ...
 - 5: return rN
-

Algorithm 2 SPN as a for loop

Inputs: **IN:** SPN leaf nodes as a vector of size **m**. **O:** Vector of size **n** with binary indicators to select between sum or product for each arithmetic operation. **B,C:** Vectors with pointers to first and second operand of each arithmetic operation.

- 1: $A \leftarrow []$ // a vector of size $m+n$
 - 2: $A[0..m-1] \leftarrow IN$ // initialize with inputs
 - 3: **for** $i = 0; i < n; i++$ **do**
 - 4: **if** $O[i] == \text{SUM}$ **then**
 - 5: $A[i+m] = A[B[i]] + A[C[i]]$
 - 6: **else**
 - 7: $A[i+m] = A[B[i]] \times A[C[i]]$
 - 8: **return** $A[m+n-1]$
-

III. SPNS ON CPU AND GPU

To evaluate the performance of general-purpose platforms, SPNs are implemented on a CPU and a GPU.

CPU: An Intel core i5-7200 CPU is used to execute SPNs as lists of operations in C (alg. 1), giving the compiler maximal freedom to reorder and parallelize operations. The for-loop implementation (alg. 2) is also evaluated, but it consistently performs slower than alg. 1

GPU: GPU offers parallelism in the form of a single instruction multiple thread (SIMT) mode. Multiple operations of SPNs can be executed in parallel by scheduling on multiple threads. However, the following two problems arise while doing so, due to the irregularity of SPN graphs.

1) Inter-thread data transfers: Inter-thread data transfers are needed when threads use data computed by other threads, which frequently happens when an SPN is scheduled across multiple threads. Such transfers can be allowed only after thread synchronization, as GPU does not guarantee lock-step execution of threads. To minimize these synchronization overheads in our experiments, SPNs are decomposed into groups of independent nodes (colored in fig. 2). As nodes in a group are not dependent on other nodes of the same group, they can be executed on any thread without synchronization. When executing a new group, threads still have to be synchronized with the CUDA primitive `__syncthreads()`. Moreover, this inter-thread communication needs the data to be in either GPU *shared* memory or *global* memory. If data size fits, the shared memory is used for low-latency inter-thread transfers, otherwise the global memory is used with L1 caching enabled.

2) Irregular memory accesses: Data transfers from the shared memory become a bottleneck if multiple threads in a

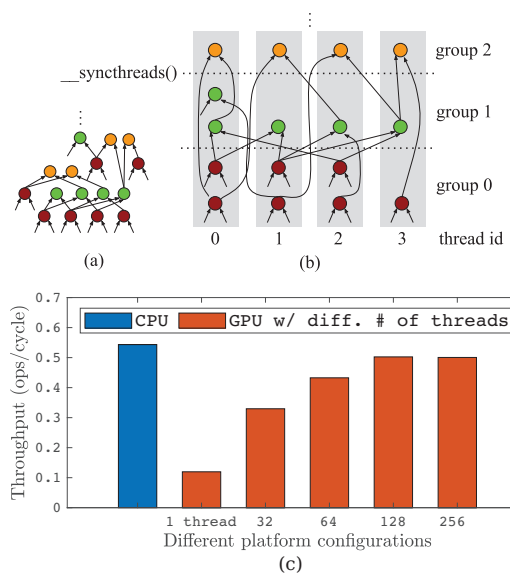


Fig. 2. Execution of an SPN on multiple GPU threads: (a) decompose SPN into 'groups' based on dependencies, and (b) execute on multiple threads and sync them before executing a new group. Throughput comparison of CPU and GPU in (c)

GPU warp access the same bank, resulting in bank conflicts. These bank conflicts are minimized with a graph coloring-based bank allocation that ensures that the threads in a warp access to different banks.

The pseudocode of the developed CUDA implementation is shown in alg. 3. Several implementations are evaluated with varying number of threads in 1 thread block, on the Nvidia Jetson TX2 embedded GPU platform. An SPN trained on a benchmark in [7] is used for performance measurement, with the resulting throughput shown in fig. 2(c). The GPU implementation utilizing a single thread expectedly performs worse than the CPU, as a GPU CUDA core is simpler than a superscalar CPU core. Use of 256 threads, however, only increases the throughput by a factor 4.1x, a sublinear scaling due to the following reasons:

- Overhead of thread synchronization [4].
- Limited shared memory bandwidth. In Jetson TX2 GPU, shared memory has 32 banks among 128 CUDA cores. As all the threads read from (and write to) the shared memory, its bandwidth becomes a bottleneck.
- Thread divergence due to selection between sum and product operations leads to inactive threads in a warp.

An efficient processor for SPNs should strive for improved parallelization by avoiding these bottlenecks.

IV. SPN PROCESSOR ARCHITECTURE

We propose a custom processor architecture (fig. 3) to alleviate the bottlenecks in SPN parallelization. It uses a lightweight method to exchange data among processing elements, and contains an appropriate memory hierarchy to handle irregular reads and writes efficiently. The important features are the following:

Algorithm 3 SPNcuda(IN, B, C, O)

```

kernel inputs: Same as inputs of alg. 2
1: int i= blockDim.x*blockIdx.x + threadIdx.x //Thread ID
2: const int t= TOTAL_THREADS, x= length(IN)
3: __shared__ A[] // a vector in shared memory
4:
5: for j = 0; j < x; j++ do
6:   A[i + j*t] = IN[i + j*t] // copy inputs to shared memory
7:   __syncthreads()
8:
9: // compute first group of nodes
10: if O[i] == SUM then
11:   A[i + x*t] = A[B[i]] + A[C[i]]
12: else
13:   A[i + x*t] = A[B[i]] × A[C[i]]
14: if O[i + t] == SUM then
15:   A[i + (x+1)*t] = A[B[i + t]] + A[C[i + t]]
16: else
17:   A[i + (x+1)*t] = A[B[i + t]] × A[C[i + t]]
18: ...
19: __syncthreads()
20: // compute second group of nodes, and so on

```

Processing element (PE): The PE is a flexible arithmetic unit that can perform all the needed arithmetic operations (+, ×). It can additionally be configured to forward either of its inputs to output without performing any operation. The output of the PE is registered to allow adequate clock frequency.

Trees of PEs: The processor's datapath consists of *trees* of PEs that enable local reuse of data, by avoiding frequent writebacks to the register file. Multiple trees allow execution of independent subgraphs of an SPN in parallel.

Register file and register writes: Each PE tree writes to a private banked register file. PEs at different levels of the tree can write to a different number of banks. As shown in fig. 3 bottom left, PEs at the first level can write to 2 respective banks, the second level can write to 4 banks, and so on.

Crossbar and register reads: The inputs of the PE trees are serviced by the register file via a crossbar, allowing PE trees to read any of the banks. Multiple inputs, however, cannot access the same bank in a given cycle. The crossbar is a combinational block with no internal memory.

Data memory: All the register files load/store data together in the form of a vector from a single address of the data memory, limiting irregular accesses to the register file, while requiring only vectorized transactions to the data memory.

Programmability: The architecture is made programmable with a custom VLIW instruction set, that can configure the trees and crossbar every clock cycle, copy data within register banks, and load/store from data memory.

Compilation: A custom compiler is developed to effectively schedule SPNs on the processor. The compiler directly takes as input the SPNs generated from tools like [5]. The compiler allocates register banks to intermediate results while trying to

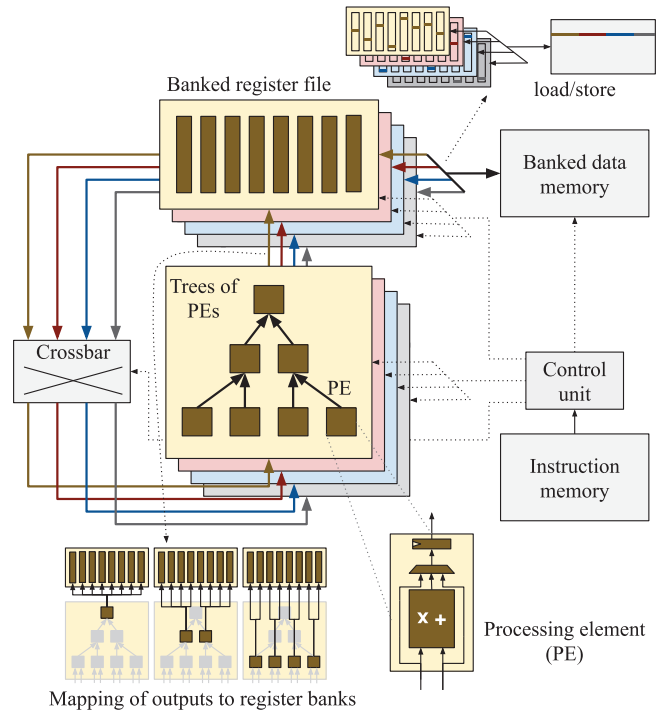


Fig. 3. The SPN processor architecture based on trees of processing elements (PE). A PE tree writes outputs to a private banked register file, while the inputs are serviced via a crossbar, allowing trees to access other register files.

minimize the register read/write bank conflicts. This allocation has to happen in tandem with the placement of operation on the PE, since PEs cannot write to all register banks. Furthermore, the compiler also reorders operations to minimize read-after-write hazards due to pipelining in PE trees. Finally, the intermediate data from the register files are spilled carefully to the data memory to minimize load-stores. In the end, the compiler generates a list of the custom VLIW instructions that can be executed directly on the processor.

In summary, the architecture allows local reuse of intermediate results by using *trees* of PEs, enables flexible movement of data with the crossbar, and efficiently supports irregular data fetches by using banked register files. The custom compiler allows to effectively use these architectural features.

V. EXPERIMENTAL RESULTS

We experimentally benchmark the throughput of the proposed processor architecture against optimized CPU and GPU implementations, with SPNs trained on a suite of standard benchmarks [3], [7] using the algorithm in [5]. The CPU and GPU implementations are as described in sec. III. For a fair comparison, the performance of the proposed processor is evaluated for configurations that consume less resources than the GPU (summarized in Table I): 32 banks in the register file, same as the amount of GPU shared memory banks, with 64 registers in each bank.

Two configurations of the custom processor with different arrangements of PEs are evaluated to highlight the impact of

the tree arrangement. The *Ptree* configuration uses 2 trees each with 4 levels of PEs (30 PEs in total). The *Pvect* configuration removes the trees and uses only the lowest level PEs (16 PEs in total). Note that both configurations are identical in every aspect (crossbar, register file etc.), except for the arrangement and the number of PEs.

Throughput for the CPU and GPU is measured by averaging the runtime for 100k SPN executions, resulting in the peak performances of 0.55 and 0.95 effective operations/cycle respectively. For accurate throughput measurement of our processor, a cycle-accurate model is developed in the MyHDL framework [2]. The resulting performance is compared in fig. 4 in terms of operations/cycle. The throughput of *Ptree* is at least 12x higher than the CPU and GPU, with a peak performance of 11.6 operations/cycle. Moreover, *Ptree* performs 2x better than *Pvect*, confirming the benefit of the tree arrangement of the PEs.

VI. RELATED WORK

In [13], authors implemented SPNs on an Nvidia 1080Ti discrete GPU using Tensorflow and observed much lower throughput than a CPU. However, Tensorflow is ill-suited for SPNs as it launches a new GPU kernel for every operation in SPN, incurring prohibitively high overhead. We show that a custom CUDA-based implementation can achieve throughput similar to a CPU, even on an embedded GPU platform.

In recent years, a few works have proposed custom hardware for SPNs acceleration. Work in [12]–[14] uses fully-parallel spatial architectures that implement every operation in SPNs as a unique hardware operator. This is typically intended for

FPGA implementation. These fully unrolled implementations, however, are not scalable to the large SPNs required for real applications. In [11], authors proposed a processor with multiple processing elements that use a common data bus to share intermediate results, making the bus a severe bottleneck. The processor in [6], proposed for a different probabilistic reasoning workload, has similarities with our work. However, it has unnecessary hardware redundancy (multiple crossbars) while having stringent requirements for the graph structure (only polytrees), making it unsuitable for SPNs. Our work retains its advantages while being more general.

VII. CONCLUSION

An efficient processor architecture is proposed to accelerate Sum-product networks, an important compute kernel for probabilistic reasoning. The processor contains hardware features suitable for the irregular graph processing of SPNs, such as trees of processing elements for data reuse, a crossbar for flexible sharing of data, and independent-addressable register banks for irregular data reads and writes. The performance of the processor is compared with CPU and an highly-optimized GPU CUDA implementation, showing 12x higher throughput while using fewer computational and memory units.

REFERENCES

- [1] M. Chavira and A. Darwiche, "On probabilistic inference by weighted model counting," vol. 172, no. 6-7. Elsevier, 2008, pp. 772–799.
- [2] J. Decaluwe, "Myhdl manual," 2019, version 0.11.
- [3] D. Dua and C. Graff, "UCI machine learning repository," 2017. [Online]. Available: <http://archive.ics.uci.edu/ml>
- [4] J. Letendre, "Understanding and modeling the synchronization cost in the gpu architecture," 2013.
- [5] Y. Liang, J. Bekker, and G. Van den Broeck, "Learning the structure of probabilistic sentential decision diagrams," in *Proceedings of the 33rd Conference on Uncertainty in Artificial Intelligence (UAI)*, 2017.
- [6] M. Lin, I. Lebedev, and J. Wawrzyniak, "High-throughput bayesian computing machine with reconfigurable hardware," in *Proceedings of the 18th annual ACM/SIGDA international symposium on Field programmable gate arrays*. ACM, 2010, pp. 73–82.
- [7] D. Lowd and J. Davis, "Learning markov network structure with decision trees," in *2010 IEEE International Conference on Data Mining*. IEEE, 2010, pp. 334–343.
- [8] R. Peharz, A. Vergari, K. Stelzner, A. Molina, X. Shao, M. Trapp, K. Kersting, and Z. Ghahramani, "Random sum-product networks: A simple and effective approach to probabilistic deep learning," in *Conference on Uncertainty in Artificial Intelligence (UAI)*, 2019.
- [9] H. Poon and P. Domingos, "Sum-product networks: A new deep architecture," in *2011 IEEE International Conference on Computer Vision Workshops (ICCV Workshops)*. IEEE, 2011, pp. 689–690.
- [10] M. Qu and J. Tang, "Probabilistic logic neural networks for reasoning," *arXiv preprint arXiv:1906.08495*, 2019.
- [11] J. Schumann, K. Y. Rozier, T. Reinbacher, O. J. Mengshoel, T. Mbaya, and C. Ippolito, "Towards real-time, on-board, hardware-supported sensor and software health management for unmanned aerial systems," *International Journal of Prognostics and Health Management*, 2015.
- [12] N. Shah, L. I. G. Olascoaga, W. Meert, and M. Verhelst, "Problp: A framework for low-precision probabilistic inference," in *Proceedings of the 56th Annual Design Automation Conference 2019*. ACM, 2019, p. 190.
- [13] L. Sommer, J. Oppermann, A. Molina, C. Binnig, K. Kersting, and A. Koch, "Automatic mapping of the sum-product network inference problem to fpga-based accelerators," in *IEEE 36th International Conference on Computer Design (ICCD)*, 2018, pp. 350–357.
- [14] S. Zermani, C. Dezan, R. Euler, and J.-P. Diguet, "Bayesian network-based framework for the design of reconfigurable health management monitors," in *2015 NASA/ESA Conference on Adaptive Hardware and Systems (AHS)*. IEEE, 2015, pp. 1–8.

TABLE I

COMPUTE AND MEMORY DETAILS OF DIFFERENT PROCESSING PLATFORMS

Platform	Compute units	Immediate memory size	memory banks
CPU	2 arith. units in a superscalar core	168 80b registers 32 KB L1 cache	16
GPU	128 CUDA cores	64K 32b registers 64 KB shared mem.	32
Ours (<i>Pvect</i>)	16 PEs	2K 32b registers 64 KB data mem.	32
Ours (<i>Ptree</i>)	30 PEs	2K 32b registers 64 KB data mem.	32

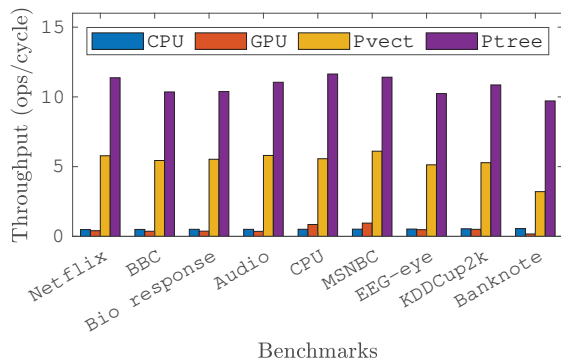


Fig. 4. Throughput comparison of different platforms on several benchmarks.