

# Low-Latency Asynchronous Logic Design for Inference at the Edge

Adrian Wheeldon, Alex Yakovlev, Rishad Shafik and Jordan Morris  
Microsystems Group, Newcastle University, Newcastle Upon Tyne, UK  
Email: adrian.wheeldon@ncl.ac.uk

**Abstract**—Modern internet of things (IoT) devices leverage machine learning inference using sensed data on-device rather than offloading them to the cloud. Commonly known as *inference at-the-edge*, this gives many benefits to the users, including personalization and security. However, such applications demand high energy efficiency and robustness. In this paper we propose a method for reduced area and power overhead of self-timed early-propagative asynchronous inference circuits, designed using the principles of learning automata. Due to natural resilience to timing as well as logic underpinning, the circuits are tolerant to variations in environment and supply voltage whilst enabling the lowest possible latency. Our method is exemplified through an inference datapath for a low power machine learning application. The circuit builds on the Tsetlin machine algorithm further enhancing its energy efficiency. Average latency of the proposed circuit is reduced by  $10\times$  compared with the synchronous implementation whilst maintaining similar area. Robustness of the proposed circuit is proven through post-synthesis simulation with 0.25 V to 1.2 V supply. Functional correctness is maintained and latency scales with gate delay as voltage is decreased.

## I. INTRODUCTION

There is an accelerating demand for connected devices in the internet of things (IoT) [1]. Such devices often comprise a sensing aspect, collecting environmental or personal data, for providing useful monitoring and decisions for transforming our everyday life. The sensors collect vast amounts of data which must be processed into a usable or more manageable form. Traditionally this was done by offloading the data into cloud compute servers, usually over a wireless medium. However this paradigm is quickly becoming unworkable as IoT devices expand well into the billions [2]. Generated data sizes become overwhelming, wireless data transmissions violate power budgets, and we see a shift towards data processing at the edge [3]. Designers of IoT products are turning to machine learning (ML) in order to extract meaningful features from the sensed data. Such products are often powered by batteries or energy harvesters which demand low power and energy efficiency, as well as robustness to supply variations [2].

There are several ML algorithms which may be suited to such applications, with neural networks (NNs) in widespread usage thanks to their often state-of-the-art accuracy and powerful hardware/software ecosystem. Hyperdimensional computing has also emerged in recent years with applications in low power systems [4]. Recently the Tsetlin machine (TM) algorithm has been proposed as a promising ML algorithm based on Tsetlin automata—specialized learning automata. The Tsetlin automata use reinforcement learning locally, together creating

an ensemble learning effect on the global scale which is used to compose logic clauses. Existing hardware based on TM offers a new direction for ML whose inference engine is based on logic with little arithmetic [5]. The logic-based underpinnings of the TM algorithm provide opportunities for low power and energy efficient ML hardware design in the IoT.

In this work we apply an asynchronous circuit design methodology [6] to the TM algorithm. By removing the pairing between clock and supply voltage as in the synchronous digital designs, it enables an aggressive voltage scaling [7] for reduced energy per inference and also adds robustness to environmental variations. Although we use ML as the key application driver, it is possible these techniques can also be applied in other application areas. Our method is built on *dual-rail* circuits with early propagation [8]. Dual-rail is an asynchronous circuit design style in the family of quasi delay insensitive (QDI) circuits. It is inherently robust to circuit delay variations which means it can operate across a wide range of supply voltages and temperatures. This usually comes at the cost of duplicated logic and completion detection (CD) overhead [9]. In our design we carefully select circuit topology to minimize such duplications. Additionally we use timing optimizations to reduce overhead from CD.

*Nomenclature:* Positive- and negative-rail signals are denoted  $x^p$  and  $x^n$  respectively.  $x_m$  denotes the  $m^{\text{th}}$  signal in the bit vector  $x$ .  $V \rightarrow S$  denotes a transition on a dual-rail signal from a valid codeword to a spacer. Vice versa for  $S \rightarrow V$ .

*Major Contributions of this paper:*

- 1) application of early-propagative, reduced-overhead self-timed dual-rail circuits to ML inference; and
- 2) analysis of operand and delay probability distributions in the ML inference circuit.

*Paper Organization:* Section II introduces the concepts of the TM algorithm. Section III first briefly introduces the principles of dual-rail circuits before describing our reduced CD scheme. Section IV presents our dual-rail inference datapath design with in-depth analyses. We finally conclude our findings in Section V.

## II. TSETLIN MACHINE OVERVIEW

The main inference component of the TM is the conjunctive clause which uses propositional logic expressions to produce a vote. The composition of each clause (determined by inclusion of literals) is controlled by the action outputs of a team of

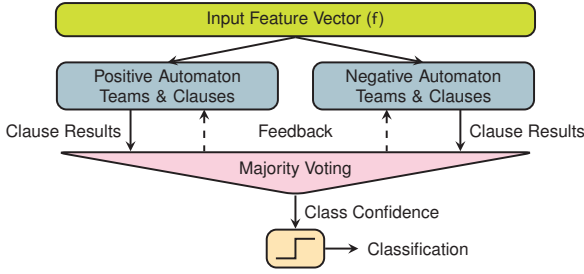


Figure 1. Simplified overview of a single Tsetlin machine (TM) classifier.

Tsetlin automata. For inference, the Tsetlin automata themselves are not required. Following a number of reinforcement steps, the automata decide whether their associated literal should be excluded from (action 1) or included in (action 2) the clause. Figure 1 illustrates a TM classifier with automaton teams and conjunctive clauses as one block for brevity.

Each clause can produce a vote for its class. Half of the clauses can vote positively, while the other half of the clauses can vote negatively. The inclusion of inhibition in the voting system enables non-linearity in the inference process. The votes are summed in a majority vote to produce a collective result which gives an indication of confidence. This confidence is used to influence future decisions of the automata [10].

A simple thresholding function can be used to generate the final classification output. If the votes are positive (or zero), the input data is determined to belong to the class. For a negative sum the input data is determined to be *not* in the class.

For purposes of studying inference, we abstract the Tsetlin automaton action outputs to the circuit's environment and concentrate on only the clauses calculation and majority voting.

### III. SELF-TIMING METHODOLOGY

In dual-rail logic two wires are used to encode a codeword. For a single bit  $x$ , the dual-rail encoding consists of the positive and negative rails  $\{x^p, x^n\}$ .  $x = 0$  is encoded as  $\{0, 1\}$ , and  $x = 1$  is encoded as  $\{1, 0\}$ . One of the remaining states— $\{0, 0\}$  or  $\{1, 1\}$ —is chosen to represent the empty state, referred to as a *spacer*, which separates valid codewords temporally so they can be distinguished from each other. Care must be taken to correctly handle spacer in the design, otherwise data hazards could occur where one valid overtakes another [11]. The remaining state is forbidden and must be avoided by design.

In our design we abide by the following requirements to ensure correct circuit operation:

- 1) Monotonic switching at the PIs.
- 2) Monotonic switching within the circuit.
- 3) Acknowledgment of  $S \rightarrow V$  on POs.
- 4)  $V \rightarrow S$  on POs and internal signals before new PIs applied.
- 5) PIs must transition  $S \rightarrow V$  and  $V \rightarrow S$  for each operand.
- 6) PIs transition  $V \rightarrow S$  only after  $S \rightarrow V$  on POs.

Requirements 1, 5 and 6 are assumed as part of the circuit's environment. To ensure Requirement 2, the circuit must be

constructed solely from unate logic gates. To maintain monotonicity we must exclude non-unate logic gates (e.g. XOR and XNOR) from our library when generating the dual-rail netlist. Requirement 3 is taken care of by CD insertion. Requirement 4 can either be assumed as part of the environment or a delay can be added to the falling edge of CD assertion. The latter will be discussed in Section III-A.

#### A. Reduced Completion Detection Scheme

CD which acknowledges both  $S \rightarrow V$  and  $V \rightarrow S$  at the primary outputs (POs) is expensive to implement due to the vast amount of complex C-elements required [12]. By indicating only  $S \rightarrow V$  transitions we can significantly reduce the overhead of CD by using a small number of simple gates.

Full CD on *internal signals* is even more costly and removes the possibility of early propagation. Its job is to ensure  $S \rightarrow V$  and  $V \rightarrow S$  occurs on internal nets for each operand. Internal CD can be safely omitted by giving a grace period for the internal signal to reset to spacer before applying new primary inputs (PIs). Codeword validity and correct operation can still be guaranteed as long as Requirements 4 to 6 are met.

In order to meet Requirement 4 there must be a sufficient grace period from application of spacer at the PIs until application of the next valid at the PIs. The grace period can be determined by using static timing analysis to find the maximum possible  $S \rightarrow V$  time on all nodes of the circuit. Consequently the grace period can be guaranteed by either 1) the circuit environment waiting for the required grace period; or 2) an appropriate delay built into the done signal of the CD. The required delay can be calculated as  $t_d = t_{int} - t_{io}$ , where  $t_{int}$  is the maximum internal net  $V \rightarrow S$  time, and  $t_{io}$  is the maximum  $V \rightarrow S$  time from the PIs to POs.  $t_{int}$  must include false paths. It is these false path which lead to the distinction between  $t_{int}$  and  $t_{io}$ . Since there may be some margin added to  $t_d$ , or due to implementation of the delay  $t_d$  may be greater than the requirement, the actual timing of the  $1 \rightarrow 0$  transition of done can be calculated by  $t_{done1 \rightarrow 0} = t_{io} + t_d$ .

### IV. INFERENCE DATAPATH

The inference datapath of the TM is derived from the full TM diagram (Figure 1). The Tsetlin automata and their feedback are not required for inference. Only the *exclude* action output is required from the Tsetlin automata teams. In the diagram this is abstracted to the PI e. We split the majority voting of the TM into two sections. Firstly we distinctly count all positive votes and negative votes by means of population counts. Secondly the two counts are compared using a magnitude comparator to determine the winner. The result of the comparison is taken as the classifier outcome.

There are several ways to construct the circuit architecture. This architecture has been chosen due to the simplicity and efficiency of the asynchronous magnitude comparator as will become clear in Section IV-C.

All PIs and POs of the circuit are dual-rail encoded. These can interface natively with other dual-rail signals, or with synchronous circuits using converters [11].

### A. Clause Calculation

The  $e$  input to the inference datapath controls whether the corresponding feature input ( $f$ ) will be excluded from a clause computation. We use OR gates to form a mask of each feature input in each clause. The partial clause values,  $pc$ , must be aggregated using an AND tree in order to evaluate the entire clause comprising input from all  $f$  and their associated automaton actions. The *exclude* signals ( $e$ ) from the Tsetlin automata mask  $f_m$  and  $\bar{f}_m$  feature inputs causing logic-1 at the AND gate inputs. If  $e_{2m}$  (resp.  $e_{2m+1}$ ) is logic-0 (ie. the feature input should be *included* in the clause calculation), the value of  $f_m$  (resp.  $\bar{f}_m$ ) is passed through to the AND gate to be evaluated. The partial clause evaluation circuit is replicated as many times as there are feature inputs to the TM.

Since  $f_m$  will be dual-rail encoded in our system, we do not need to generate  $\bar{f}_m$  internally. By performing direct mapping of a single-rail circuit, and along with negative gate optimization [11], we arrive at the optimized dual-rail circuit in Figure 2. All signal paths in this circuit have a single inversion—satisfying spacer requirements and giving the block an *inverting spacer* overall.

### B. Population Count

We base our population count circuit on the optimized design of Dalalah [13]. The eight-input design comprises nine half-adders, two full-adders and two OR gates and is illustrated in Figure 2. Each wire in the diagram represents two signals which form the dual-rail encoding. The dual-rail OR gate is internally constructed from one OR gate and one AND gate. The dual-rail half-adders are constructed using two complex gates and two simple gates each. There is no spacer inversion within the half-adders as all signal paths have an even number of inversions. The dual-rail full-adder is constructed from six complex gates, two simple gates and four inverters [6]. It has inverted spacers on carry-in and carry-out with respect to the other inputs and outputs, therefore we must accommodate for these in the population count design by adding spacer inverters: 1) between  $HA_8$  and  $FA_0$ ; 2) between  $FA_1$  and the  $y_3$  output. The resulting dual-rail population count circuit has no spacer inversion overall, therefore the output spacer will have the same polarity as the input spacer.

### C. Magnitude Comparator

The magnitude comparator compares the number of votes from the positive and negative Tsetlin automaton teams. A larger number of positive votes indicates that the input pattern belongs to the class in question, and conversely, a larger number of negative votes indicates that the input pattern *does not* belong to the class in question.

The magnitude comparator is based on a request architecture [6] and compares the operands in bit-pairs, starting from the most significant bit. Once a difference is found, the answer is known, and the remainder of the bits need not be compared. This architecture enables huge average-case latency improvement over a synchronous counterpart. Energy savings

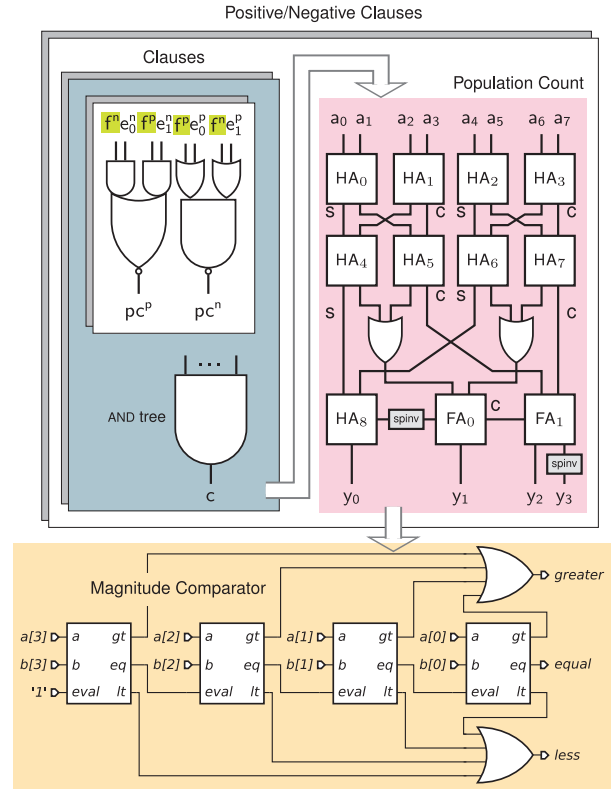


Figure 2. Block diagram of the Tsetlin machine (TM) inference path. Including dual-rail circuits for partial clause evaluation; and population count architecture, where each wire represents two signals with dual-rail encoding. spinv: spacer inverter.

are also made by due to saved switching power on the lower bits when the operands differ by a large magnitude.

Since the comparator's outputs (less, equal, and greater) are mutually exclusive, we take advantage of this in the asynchronous design. We use a 1-of-3 encoding on the output instead of the usual dual-rail—1-of-n encoding being a superset of dual-rail. Provided a spacer separates the valids, the switching of 1-of-n codes is monotonic [14], therefore satisfying Requirement 2. Without this trick, three sets of dual-rail signals would be required at the comparator outputs at the expense of more logic to drive these signals. The inputs to the comparator are dual-rail encoded.

### D. Inference Datapath Results

The inference datapath was synthesized using SYNOPSIS DESIGN COMPILER for two different 65 nm silicon libraries. UMC LL is a commercially available, low-leakage library which we use with nominal 1.2 V supply and TT corner. FULL DIFFUSION is a custom library aimed at high performance subthreshold operation [15]. It uses a full diffusion sizing strategy with non-minimum-length transistors in order to mitigate subthreshold effects. For this silicon library the circuit is first synthesized at TT corner for nominal 1.2 V supply and results are shown with supply voltage in the range 0.25 V to 1.2 V.

Table I  
COMPARISON OF SINGLE-RAIL AND DUAL-RAIL CIRCUITS AFTER SYNTHESIS.

Technology	Design	Cell Area	Sequential Area	Avg.Power ( $\mu$ W)	Leakage Power (nW)	Avg.Latency (ps)	Max Latency (ps)	$t_{v \rightarrow s}$ (ps)	Avg.Inferences (Millions $s^{-1}$ )
UMC LL	Single-rail	1800	1300	470	75	2100	2100	—	480
	Proposed Dual-rail	2000	1100	660	73	260	3100	3100	300
FULL DIFFUSION	Single-rail	3400	2500	990	37	2400	2400	—	410
	Proposed Dual-rail	3800	2400	1700	62	220	1900	1700	510




Figure 3. Scaling of dual-rail datapath latency with supply voltage for the FULL DIFFUSION library.

Results in Table I show similar cell areas for both single-rail and dual-rail designs for each silicon library. This is possible due to the careful choice of dual-rail circuit architecture and the reduced completion detection scheme. The dual-rail clause computation and magnitude comparator are more area efficient than their single-rail counterparts due to exploitation of dual-rail encoding and clever use of 1-of-3 encoding respectively.

For the area of the sequential cells we count flip-flop area for the single-rail designs and C-element area for dual-rail designs. The sequential area is similar between designs, despite the dual-rail design having twice as many sequential cells due to the doubled input rails. The dual-rail circuit uses C-elements as latches. These comprise four simple gates in the FULL DIFFUSION library (due to lack of AOI32 cells) and a single complex gate in the UMC LL library. Note that the cell area varies dramatically between the libraries due to transistor sizing—UMC LL being minimally-sized for superthreshold and FULL DIFFUSION larger for subthreshold operation. The number of cells does not vary significantly.

Latency is measured from  $s \rightarrow v$  in the dual-rail designs, and the clock period defines the latency for single-rail designs. The dual-rail circuit enables  $10\times$  reduction in average latency thanks to early propagation. Average throughput is worsened however, due to the lengthened logic path and the need for the additional  $v \rightarrow s$  transition. Although the dual-rail switching power is greater due to higher inherent activity factor, the computation energy is reduced due to increased throughput.

Throughput period is defined by the single-rail circuit's clock period. For the dual-rail design, throughput period is determined by  $t_{s \rightarrow v} + t_{s \rightarrow v}$  so that the PIs are ready for the next operand.  $t_{v \rightarrow s}$  has the same magnitude as  $\max(t_{s \rightarrow v})$ .

Figure 3 shows the effects of supply voltage on datapath latency. The latency increases exponentially as the supply voltage is reduced from 0.6 V to 0.25 V. The key point is that the circuit functionality is guaranteed across the whole supply voltage range thanks to the requirements in Section III and without any alteration to the hardware.

## V. CONCLUSION

In this paper we have demonstrated an asynchronous, self-timed inference datapath design with area and power of equal orders of magnitude to the synchronous equivalent. Early propagation enables  $10\times$  lower inference latency than the equivalent synchronous circuit on average. The savings are enabled by a reduced CD scheme which can be applied to any dual-rail asynchronous circuit. The new scheme introduces a timing assumption which can be incorporated into the CD circuit, so that the circuit environment does not need to be adapted. This type of low-latency circuit can have applications in speech recognition for wearables and other low-power applications where inference latency is of particular importance.

In future work we will apply asynchronous design styles to the training datapath of the TM algorithm in order to enable a fully-asynchronous ML hardware capable of on-chip learning.

## REFERENCES

- [1] M. Capra *et al.*, "Edge computing: A survey on the hardware requirements in the Internet of Things world," p. 100, Apr. 2019.
- [2] R. Shafik, A. Yakovlev, and S. Das, "Real-power computing," *IEEE Trans. Comput.*, vol. 67, no. 10, pp. 1445–1461, 2018.
- [3] "Arm AI Platform Solutions Brief," Arm Limited, Tech. Rep., 2020.
- [4] A. Burrello *et al.*, "Hyperdimensional Computing with Local Binary Patterns: One-Shot Learning of Seizure Onset and Identification of Ictogenic Brain Regions Using Short-Time iEEG Recordings," *IEEE Trans. Biomed. Eng.*, vol. 67, no. 2, pp. 601–613, Feb. 2020.
- [5] A. Wheelodon *et al.*, "Learning automata based energy-efficient AI hardware design for IoT applications," *Philos. Trans. R. Soc. A Math. Phys. Eng. Sci.*, vol. 378, no. 2182, 2020.
- [6] —, "Self-timed, minimum latency circuits for the internet of things," *Integration*, vol. 69, pp. 138–146, Nov. 2019.
- [7] R. Diamant, R. Ginosar, and C. Sotiriou, "Asynchronous sub-threshold ultra-low power processor," in *2015 25th Int. Work. Power Timing Model. Optim. Simul.* IEEE, Sep. 2015, pp. 89–96.
- [8] C. F. Brey and J. D. Garside, "Early Output Logic using Anti-Tokens," in *Proc. IEEE/ACM Int. Conf. Comput. Des.*, 2006, pp. 158–163.
- [9] A. Yakovlev, P. Vivet, and M. Renaudin, "Advances in Asynchronous logic: from Principles to GALS & NoC, Recent Industry Applications, and Commercial CAD tools," in *Proc. Conf. Des. Autom. Test Eur.*, 2013.
- [10] O.-C. Granmo, "The Tsetlin Machine - A Game Theoretic Bandit Driven Approach to Optimal Pattern Recognition with Propositional Logic," Apr. 2018.
- [11] D. Sokolov, "Automated synthesis of asynchronous circuits using direct mapping for control and data paths," Ph.D. dissertation, Newcastle University, 2006.
- [12] J. Sparsø and S. Furber, *Principles of Asynchronous Design: A Systems Perspective*. Kluwer Academic Publishers, 2001.
- [13] A. Dalalah, S. Baba, and A. Tubaihat, "New Hardware Architecture for Bit-Counting," in *Proc. 5th WSEAS Int. Conf. Appl. Comput. Sci.*, 2006, pp. 118–128.
- [14] W. Bainbridge *et al.*, "Delay-insensitive, point-to-point interconnect using m-of-n codes," in *Ninth Int. Symp. Asynchronous Circuits Syst. 2003. Proceedings*. IEEE Comput. Soc, 2003, pp. 132–140.
- [15] J. Morris *et al.*, "Unconventional Layout Techniques for a High Performance, Low Variability Subthreshold Standard Cell Library," in *2017 IEEE Comput. Soc. Annu. Symp. VLSI*. IEEE, Jul. 2017, pp. 19–24.