POST-DIGITAL - European Training Network on Post-Digital Computing [GA860360]

Document Details

Title	Deliverable 4.3. Report on comparison of neuromorphic computation solutions, with respect to performance on standard tasks, speed, energy consumption and footprint
Deliverable number	D4.3
Deliverable Type	Report (public)
Deliverable title	Report on comparison of neuromorphic computation solutions, with respect to performance on standard tasks, speed, energy consumption and footprint
Work Package	WP4- Benchmarking and applications
Description	Report on comparison of neuromorphic computation solutions, with respect to performance on standard tasks, speed, energy consumption and footprint
Deliverable due date	30/11/2023
Actual date of submission	27/11/2023
Lead beneficiary	ULB
Version number	V1.0
Status	Final

Dissemination level

PU	Public	X
CO	Confidential, only for members of the consortium (including Commission Services	

Project Details

Grant Agreement	860360
Project Acronym	POST-DIGITAL
Project Title	POST-DIGITAL - European Training Network on Post-Digital Computing
Call Identifier	H2020-MSCA-ITN-2019
Project Website	https://postdigital.astonphotonics.uk/
Start of the Project	1 April 2020
Project Duration	48 months

This project has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 860360.

EXECUTIVE SUMMARY

This deliverable presents a comparison of neuromorphic computation solutions, with respect to performance on standard tasks, speed, energy consumption and footprint. It has been completed with the contribution of all POST-DIGITAL ESRs under the lead of Professor Serge Massar from Université Libre de Bruxelles.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF ACRONYMS

ADC	Analog-to-Digital Converter
AiPT	Aston Institute of Photonic Technologies
BER	Bit-error rate
DAC	Digital-to-Analog Converter
ELM	Extreme learning machines
ESR	Early Stage Researcher
ETN	European Training Network
FPGA	Field Programmable Gate Array
NMSE	Normalised Means Square Error
NRMSE	Normalised Root Means Square Error
POST-DIGITAL	Project ‘POST-DIGITAL - European Training Network on Post-Digital Computing’ EC GA 860360
RC	Reservoir Computer
SER	Symbol error rate
TPA	Two-photon absorption
WER	Word error rate

1 INTRODUCTION

We compare different neuromorphic computation solutions with respect to the type of algorithm implemented, the number of nodes (or “neurons”), the performance on some standard tasks, speed, estimates of energy consumption, and footprint.

Given the very wide variety and number of systems reported in the literature, we have focused on systems developed or used by members of POST-DIGITAL (with one exception: we have included data on a related system published by another group, see system 3). We report on 9 systems. A table summarizing the comparison can be found below, followed by detailed information about each system.

2 COMPARISON OF PERFORMANCE OF DIFFERENT NEUROMORPHIC SYSTEMS

Concerning the type of algorithms implemented, systems 1 to 5 implement the reservoir computing algorithm (system 4 is a preliminary numerical study of a system that is not yet implemented); 6 is an extreme learning machine; 7 is an accelerator for vector-matrix multiplication; 8 an event based camera supplemented by offline processing using a spiking neural network (Loihi); and 9 concerns the software development of a novel training algorithm.

Concerning the number of nodes or “neurons”, the POST-DIGITAL systems range from a dozen to 10^3.

Concerning tasks, each kind of algorithm addresses the processing of different kinds of information. Therefore, the tasks on which different algorithms are tested are different. If we focus on the 5 systems implementing reservoir computing, then the same task has been tackled by different systems, and performance can be compared. Note however that some results are reported in terms of Normalised Means Square Error (NMSE) and some results are reported in terms of Normalised Root Means Square Error (NRMSE)). Because of the diversity of results reported, we do not present these results in the table and refer to the detailed description.

Concerning speed, the fastest implemented systems are 1 and 3 which are reservoir computers that can process one input every 300 ps, and the vector matrix multiplication accelerator that takes approximately 1 ns to process an input. The slowest system is 6 which processes one input every 0,5 s (due to the very slow refresh rate of the spectral filter used in the experiment).

Concerning energy consumption, a wide variety of results are reported. Many of the systems described are optical, and one needs to separate the energy consumption of the purely optical systems from the energy consumption of the drivers and current sources, and of the supporting electronics. Because these are proof of principle demonstrators, no real effort has been put in minimizing the overall energy consumption.

Finally concerning footprint, a big difference can be made between systems using integrated optics in which case the footprint of the optical system is on the order of 1 cm^2, and tabletop experiments in which case the footprint is a fraction of a m^2. Supporting electronics have not been packaged, and at present have a footprint of order m^2.

The present report is quite preliminary. In particular, it does not go into the potentialities of each system if it were optimized. It would thus be interesting for each system to reflect on the ultimate limits in terms of speed; of energy consumption (how much can the energy of the drivers and supporting electronics be reduced; do we need an ADC and a DAC, or can the signals remain in the analogue domain); of footprint (including all supporting electronics). Future analysis could focus on the potential of each system in terms of speed, footprint, energy consumption, etc… if the system were optimized further.

Table 1 Comparison of performance of 9 different neuromorphic systems. See detailed information in the main text.

	System	Algorithm	# neurons	Input processing time	Energy consumption	Footprint
1	High speed Opto-Electronic Delay System	Reservoir Computing	30-2000	0.3-20 ns	1W + electronics	-
2	Opto-Electronic Delay System + FPGA	Reservoir Computing	50-600	2-200 μs	0,14W (optical)+ 12W (electronics)	30*100 cm^2
3	Integrated Opto-Electronic with feedback laser	Reservoir Computing	0-250	254 ps	-	+- 1 cm^2 (for chip)
4	Evanescent coupling between resonators (numerical simulation)	Reservoir Computing	12-48	40 ps	-	< 1 cm^2 (for chip)
5	Optoelectronic using Wavelength Multiplexing (RC)	Reservoir Computing	20-40	5ns	1200W	60*80 cm^2
6	Optoelectronic using Wavelength Multiplexing (ELM)	Extreme Learning Machine	20-60	0.5 s	1600W	100*50 cm^2
7	Vector Matrix Multiplication using integrated photonics	Vector Matrix Multiplication	1000*1000	1 ns	20W	1.6*1.6 cm^2
8	Optoelectronic event-based neuromorphic system + offline spiking neurons	Feed-forward neural network	500-1000	1-100 μs	-	50*10 cm^2
9	conceptor-based regularization on desktop computer	Backpropagation through time using Conceptor-based regularization	512	(total training time 4 min)	35 W (during training)	-

2.1 HIGH SPEED OPTOELECTRONIC DELAY DYNAMICAL SYSTEM

Algorithm: reservoir computing

Output layer: digital

Number of nodes: 32 - 2094 time-multiplexed neurons

Performance on standard task

Mackey-Glass 9-step-ahead prediction: NRMSE=0.01

Santa-Fe one step ahead prediction: NRMSE=0.06

Fiber transmission equalization: BER<0.001

Speed

time per node: 11.7 ps

input processing speed: up to 2.67 GHz

E**stimated energy consumption:** 1.07W (includes laser operation, optical encoding and photodetection / does not include digital post-processing)

Footprint: table top system

Reference

Goldmann, M., Estebanez, I., Vlieg, E. A., Mirasso, C. R., Fischer, I., Argyris, A., & Soriano, M. C. (2023, June). Speeding up a time-delay photonic reservoir. In The European Conference on Lasers and Electro-Optics (p. jsiii_2_2). Optica Publishing Group.

2.2 OPTOELECTRONIC DELAY DYNAMICAL SYSTEM WITH FPGA

Algorithm: Reservoir Computing

Digital interface: FPGA

Number of “neurons” : 50 to 600

Performance on standard task

NARMA10: NMSE=0.45 (with 50 Neurons)

10 step ahead Mackey Glass prediction: NMSE=0.15 (with 50 Neurons)

Nonlinear channel equalization (noiseless): SER=0 (with 50 neurons)

Spoken Digits Recognition: WER=0.13 (with 100 neurons)

KTH Human Action Recognition Dataset: Accuracy = 90.87% (with 600 Neurons)

Speed

Time to process one neuron: 39 ns to 156 ns

Time to process one input: number of neurons * time to process one neuron. From 2 ms to 187ms.

Energy consumption

Energy consumption of photonic devices: 0.14 W

Energy consumption of supporting electronics, including FPGA board and ADC+DAC board: 12W.

Footprint
approx. 30cm*100cm

References

[1] Picco, Enrico, and Serge Massar. "Real-Time Photonic Deep Reservoir Computing for Speech Recognition." 2023 International Joint Conference on Neural Networks (IJCNN). IEEE, 2023.
[2] Picco, Enrico, Piotr Antonik, and Serge Massar. "High speed human action recognition using a photonic reservoir computer." Neural Networks (2023).

2.3 Integrated optical delay dynamical system based on a distributed-feedback laser

Algorithm: Reservoir Computing
Number of "neurons" : up to 250
Performance on standard task
Santa-Fe one step ahead prediction: NRMSE=0.109
Nonlinear channel equalization: SER=0.03 (with 120 neurons)
Speed
Time for processing of one input: 254 ps
Energy consumption: No estimate given
Footprint: approx. 1cm^2 (for the photonics integrated circuit)

Reference

Takano, K., Sugano, C., Inubushi, M., Yoshimura, K., Sunada, S., Kanno, K., & Uchida, A. (2018). Compact reservoir computing with a photonic integrated circuit. *Optics express, 26(22),* 29424-29439.

2.4 All-optical integrated dynamical system based on evanescent coupling between resonators (numerical simulation)

Algorithm: Reservoir Computing (space-multiplexed)
Number of "neurons": 12 – 48
Performance on standard task
Mackey Glass 3-step ahead prediction: NRMSE=0.037 (with Two Photon Absorption); NRMSE=0.07 (TPA + Free Carrier Dispersion); NRMSE=0.087 (TPA + Free Carrier Dispersion + Kerr effect). (Simulations).
Speed
Limited by intrinsic nonlinearity, so depends on material and resonator type, about 25 GHz for GaAs microrings
Energy consumption
Without supporting electronics: approx. 1 mW / neuron for GaAs microrings with a Q-factor of 10^5.

Footprint

A few square mm for optical chip

References

Boikov, I., Brunner, D., & De Rossi, A. (2023). Evanescent coupling of nonlinear integrated cavities for all-optical reservoir computing. *New Journal of Physics.* **25** 093056

2.5 Optoelectronic wavelength-multiplexing dynamical system (RC)

Algorithm: (Deep) Reservoir Computing

Number of “neurons” : 20 to 40

Performance on standard task

Nonlinear channel equalization (noiseless): SER=10^{-3} to 10^{-5}

Santa-Fe one step ahead prediction: NMSE: 0.04 to 0.004

Speed

Time for processing of one input: 5 ns

Energy consumption: 2500W (most energy consumption comes from the Erbium Doped Fiber Amplifiers).

Footprint: approx. 80*60 cm^2

Reference

L. Butschek, et al. Photonic reservoir computer based on frequency multiplexing. Optics Letters, 2022, 47.4: 782-785.

A. Lupo, et al. Fully analog photonic deep Reservoir Computer based on frequency multiplexing. arXiv preprint arXiv:2305.08892, 2023. (Accepted for publication on Optica)

2.6 Optoelectronic wavelength-multiplexing dynamical system (ELM)

Algorithm: Extreme Learning Machine

Number of “neurons”: 20 - 60

Performance on standard task

Nonlinear Channel Equalization: SER<10^{-3}

Mushroom Classification: 95.1%

Speed

Time for processing of one input: 0.5 s

Energy consumption

approximately 1600 W (including supporting electronics)

Footprint

100cm x 50cm

References

A. Lupo, L. Butschek, S. Massar, Photonic extreme learning machine based on frequency multiplexing. Optics express, 2021, 29.18: 28257-28276

A. Lupo, S. Massar, Parallel extreme learning machines based on frequency multiplexing. Applied Sciences, 2021, 12.1: 214.

2.7 Vector-Matrix Multiplication using Photorefractive holographic diffraction matrix in integrated photonics

Algorithm: Matrix-vector multiplication

Number of "neurons": 1000x1000 matrix

Performance on standard task

4-bit resolution with 1% BER

Speed

1GHz processing rate

10^{15} MAC/s

Energy consumption

Energy consumption including supporting electronics: 20W

Footprint

1.6x1.6 cm^2 = 250 mm^2

Reference

E. A. Vlieg, L. Talandier, R. Dangel, F. Horst, J. B. Offrein, An Integrated Photorefractive Analog Matrix-Vector Multiplier for Machine Learning. *Applied Sciences*. 2022; 12(9):4226 (concept)

2.8 Optoelectronic neuromorphic system using event based camera

Algorithm: Spiking feed-forward neural network using data recorded by an event-based camera

Number of neurons: The optical side is used as an extreme learning machine and number of neurons is not given. On the electronic side, either one or two layer with 512 neurons per layer.

Performance on standard task

The setup was task-specific for flow cytometry (classification of two particles) and only evaluated on the data measurement collected for this task.

Speed

The optical part of the setup runs in real-time, with a temporal resolution of 1µs.

The electronic system (Intel Loihi 2) is accessible only through the cloud and processes events accumulated over 10µs or 100µs.

Energy consumption

Energy consumption of the electronic system: 0.55W

Energy consumption of the optical system: unknown (simple laser)

Footprint

approx. 50cm*10cm*10cm (free space optical setup)

References

1) Training a spiking neural network on an event-based label-free flow cytometry dataset. M Gouda, S Abreu, A Lugnan, P Bienstman. arXiv preprint arXiv:2303.10632, 2023.

2) Flow Cytometry With Event-Based Vision and Spiking Neuromorphic Hardware. Steven Abreu, Muhammed Gouda, Alessio Lugnan, Peter Bienstman; Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops, 2023, pp. 4139-4147

2.9 Interpolating motion pattern using conceptor-based regularization

Algorithm: Backpropagation through time using conceptor-based regularization. Software implementation in python (using jax).

Number of neurons: 512

Performance on standard task

dataset: Motion Capture (input dim.=94,output dim.=94)

CMU_016_15 (walking)

CMU_016_55 (running)

NMSE: 0.1065

Generalization Ability: Two-shot learning of continuous interpolation between walking and running

Total time for training: ~4 Min

Processor used: AMD Ryzen 7 6800 HS Creator Edition (~4.7GHz)

Estimated energy consumption: 35W (CPU only) over ~240 secs