Third International Workshop on Theoretical and Experimental Material Computing (TEMC 2021)

Espoo, Finland
October 2021

TEMC 2021 is sponsored by the
EPRSC MARCH project EP/V006029/1 and EP/V006339/1,
the University of York, UK, and the University of Sheffield, UK

The **Third International Workshop on Theoretical and Experimental Material Computing** (TEMC 2021) is being held in Espoo, Finland, as a satellite workshop of the International Conference on Unconventional Computation and Natural Computation (UCNC 2019), 18-22 October 2021.

Material computing exploits unconventional physical substrates and/or unconventional computational models to perform physical computation in a non-silicon and/or non-Turing paradigm.

TEMC 2021 encompasses a range of theoretical and experimental approaches to material computing. The aim of the workshop is to bring together researchers from a range of connected fields, to inform of latest findings, to engage across the disciplines, to transfer discoveries and concepts from one field to another, and to inspire new collaborations and new ideas.

Programme and Organising Committee

Susan Stepney, Dan Allwood
David Griffin, Mohammad Musameh, Charles Swindells
Simon O'Keefe, Martin Trefzer
Tom Hayward, Eleni Vasilaki

Programme (Mon 18 October 2021)

14:45-15:00 EEST (UTC+3; BST+2)

Susan Stepney (chair)

Introduction and Welcome

15:00-16:00

Damien Querlioz (keynote speaker).

Equilibrium Propagation: a Road for Physics-Based Learning

16:00-16:30

coffee break

16:30-17:00

Nicholas Chancellor, Viv Kendon.

Experimental test of search range in quantum annealing

17:00-17:30

Alexander McDonnell, Matthew Dale, Martin A. Trefzer.

Delay-Feedback Reservoir Computing using a Field Programmable Analogue Array (FPAA)

17:30-18:00

Arthur Penty, Gunnar Tufte.

Sculpting the Spin Ice for Computation

18:00

Close

Equilibrium Propagation: a Road for Physics-Based Learning

Damien Querlioz

Université Paris-Saclay, CNRS, C2N, Palaiseau, France

Neuromorphic computing takes inspiration from the brain to create highly energy-efficient hardware for information processing, capable of sophisticated tasks. The resulting systems are most often preprogrammed: training neuromorphic systems on-chip to perform new tasks remains a formidable challenge. The flagship algorithm for training neural networks, backpropagation, is indeed not hardware-friendly. It requires a mathematical procedure to compute gradients, external memories to store them, and an external dedicated circuit to change the neural network parameters according to these gradients. The brain, by contrast, does not learn this way. It learns intrinsically, and its synapses evolve directly through the spikes applied by the neurons they connect, using their biophysics. This technique is very advantageous in terms of energy efficiency and device density. In this talk, I will introduce our approach towards reproducing this brain strategy of intrinsic learning exploiting device physics. I will show through simulations how we take advantage of the physical roots of an algorithm called Equilibrium Propagation (1) to design dynamical circuits that learn intrinsically with high accuracy (2–4).

1. B. Scellier, Y. Bengio, *Front. Comput. Neurosci*. 11 (2017).
2. M. Ernoult, J. Grollier, D. Querlioz, Y. Bengio, B. Scellier, *Proc. NeurIPS*, pp. 7081 (2019).
3. A. Laborieux et al., *Front. Neurosci*. 15 (2021).
4. E. Martin et al., *iScience*. 24 (2021)

---

bio: Damien Querlioz is a CNRS Researcher at the Centre de Nanosciences et de Nanotechnologies of Université Paris-Saclay. His research focuses on novel usages of emerging non-volatile memory and other nanodevices, in particular relying on inspirations from biology and machine learning. He received his predoctoral education at Ecole Normale Supérieure, Paris and his PhD from Université Paris-Sud in 2009. Before his appointment at CNRS, he was a Postdoctoral Scholar at Stanford University and at the Commissariat a l'Energie Atomique. Damien Querlioz is the coordinator of the interdisciplinary INTEGNANO research group, with colleagues working on all aspects of nanodevice physics and technology, from materials to systems. He is a member of the bureau of the French Biocomp research network. In 2016, he was the recipient of an ERC Starting Grant to develop the concept of natively intelligent memory. In 2017, he received the CNRS Bronze medal. He has also been a co-recipient of the 2017 IEEE Guillemin-Cauer Best Paper Award and of the 2018 IEEE Biomedical Circuits and Systems Best Paper Award.

Experimental test of search range in quantum annealing

Nicholas Chancellor[1] and Viv Kendon[1]

[1] *Department of Physics; Joint Quantum Centre (JQC) Durham-Newcastle, Durham University, South Road, Durham, DH1 3LE, UK*

(Dated: 12th August 2021)

Abstract for TEMC 2021. This work is recently published as Phys. Rev. A 104, 012604 (2021) DOI:10.1103/PhysRevA.104.012604. – a brief summary is given here.

We construct an Ising Hamiltonian with an engineered energy landscape such that it has a local energy minimum which is near the true global minimum solution and further away from a false minimum. Using a reverse annealing technique established in previous experiments, we design our experiment such that (at least on timescales relevant to our study) the false minimum is reached preferentially in forward annealing due to high levels of quantum fluctuations. This allows us to demonstrate the key principle of reverse annealing, that the solution space can be searched locally, preferentially finding nearby solutions, even in the presence of a false minimum. The techniques used here are distinct from previously used experimental techniques and allow us to probe the fundamental search range of the device in an alternative way.

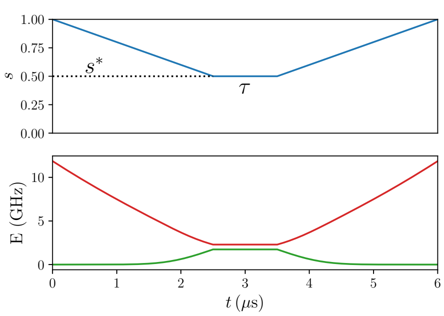

Figure 1. (a) Plot of s versus time for a reverse annealing protocol in which the device is annealed at the maximum allowed rate (which would take 5 μs to traverse from $s = 0$ to $s = 1$), with a hold time τ of 1 μs and $s^* = 0.5$. (b) The A and B energy scales of this protocol performed on the low noise QPU model. The red (dark gray) curve starting from around 10 GHz is B and the green (light gray) curve starting at around 0 GHz is A.

We perform these experiments on two flux qubit quantum annealers, one with higher noise levels than the other. We find evidence that the lower noise device is more likely to find the more distant energy minimum (the false minimum in this case), suggesting that reducing noise fundamentally increases the range over which flux qubit quantum annealers are able to search. Our work explains why reducing the noise leads to improved performance on these quantum annealers. This supports the idea that these devices may be able to search over broad regions of the solution space quickly, one of the core reasons why quantum annealers are viewed as a potential avenue for a quantum computational advantage.

The experiments use a reverse annealing protocol, rather than a standard quantum anneal, see figure 1, where the parameters A and B refer to quantum evolution under the Hamiltonian

$$\hat{H}(t) = A(t)\hat{H}_0 + B(t)\hat{H}_p.$$

The two parts of the Hamiltonian $\hat{H}_0$ provides dynamics by rotating the qubits, while $\hat{H}_p$ is the Ising Hamiltonian we are trying to find the ground state of, shown schematically in figure 2.

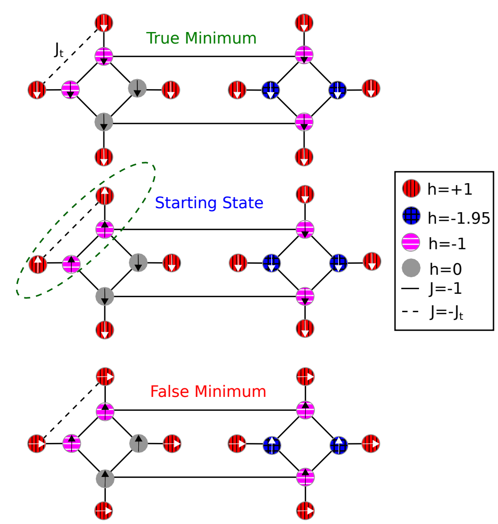

Figure 2. Hamiltonian used in the experiment and the three relevant sets of states: (a) the true minimum, (b) the starting state, and (c) the manifold of states which comprise the false minimum. Solid edges represent ferromagnetic couplings of unit strength and the dashed edge represents a tunable ferromagnetic coupling. The different coloured circles represent qubits with different fields. Arrows indicate the state of the qubit, with up (down) arrows representing $|0\rangle$ ($|1\rangle$) and horizontal arrows representing superpositions of computational basis states which may be measured as either $|0\rangle$ or $|1\rangle$. The dashed oval is a guide to the eye to show which qubits are flipped between the starting state and the true minimum-energy state.

Delay-Feedback Reservoir Computing using a Field Programmable Analogue Array (FPAA)

Alexander McDonnell[1], Matthew Dale[2], and Martin A. Trefzer[1]

[1]Department of Electronic Engineering, University of York, United Kingdom
[2]Department of Computer Science, University of York, United Kingdom

***Abstract*—Here we outline an early implementation of a Delay-Feedback Reservoir Computer on a Field Programmable Analogue Array (FPAA). The FPAA allows integration of real-time analogue signal processing with traditional digital circuitry and provides the flexibility to pre-process input data directly on the chip, on-line reconfigurability of the reservoir, and extended scalability by chaining devices together.**

I. Introduction

As Moore's law becomes unsustainable, new computing concepts and hardware are needed to solve increasingly complex and demanding computational problems. One possibility is to look towards biological systems for inspiration, such as Artificial Neural Networks. Recurrent Neural Networks (RNN), in particular, show great promise due to some unique properties, including: a non-linear response to input stimuli, high-dimensionality, and a fading memory of previous inputs. Unfortunately, RNNs are difficult to train due to the vanishing and exploding gradient problem [1]. A solution is to exploit the dynamics of randomly connected RNNs (Echo State Networks (ESN) [2]) as a "black box" and train only the outputs of the neurons through a "readout layer". ESNs have been found to perform exceptional well even with minimal training. The concept of training "black boxes" later morphed into exploiting different input-driven dynamical systems, commonly called "reservoirs", leading to the field of Reservoir Computing (RC) [3]. Various physical systems have been applied to the RC framework, even a bucket of water (see review [4]).

II. Delay-Feedback Reservoir

A type of RC of particular interest is the delay-feedback reservoir, where a single non-linear node and a delay line can be used to generate the complexity of a virtual network when the input data is time-multiplexed [5]. Consisting only of a single node and a delay, it boasts an attractive alternative compared to typical larger physical systems, however, the downside is that extensive pre-processing is needed on the input and output data.

The delay-feedback reservoir consists of three main parts: a pre-processing stage that time-multiplexes the input data with a higher frequency masking signal; the reservoir network that contains a single non-linear node and delay line; and a readout layer that demultiplexes the output of the node and contains the trained weights of the system. The number of virtual nodes is defined by the relationship between the time period of the input signal T and the time period of the masking signal θ. Typically, the more virtual nodes, the greater the memory of the system.

Several physical implementations of delay-feedback reservoirs exist, these are typically done using digital or opto-electronic techniques [4], [6]; this has led to digital-analogue hybrids that use digital processing with an analogue non-linear node. However, it is possible to utilise both digital and analogue technology on a single chip using the Field Programmable Analogue Array (FPAA).

III. Field Programmable Analogue Array (FPAA)

The FPAA allows integration of real-time analogue signal processing with traditional digital circuitry. To realise this, the FPAA utilises switched capacitor technology allowing for configurable analogue blocks (CABs) to dynamically implement configurable analogue modules (CAMs) with a wide range of functionality. Like its digital counterpart, the Field Programmable Gate Array (FPGA), the FPAA is fully reconfigurable; some models even allow for on-line reconfiguration without any data loss.

A. Implementation

The proposed delay-feedback reservoir is implemented on a Anadigm QuadApex development board which contains four FPAA chips [7]. Figure 1 shows the proposed circuit using the provided CAMs.

The heart of the circuit is the user-defined voltage transfer function CAM. This CAM allows for any arbitrary function to be programmed within an 8-bit resolution. Here, the Mackey-Glass delay differential equation is implemented as it exhibits many of the desired characteristics for RC. It is also possible to cascade functions together to create more complex RC nodes as well. After the non-linear node, the output signal is integrated with a leaky decay rate. The integration constant, specified within the CAM, determines the connectivity between the virtual nodes.

The next block is the delay CAM. This digitises the input values to an 8-bit resolution and stores them within a lookup table, and outputs the reconstructed analogue signal after a particular time. Finally, a summation CAM sums the input signal and delayed feedback.

The FPAA also features CAMs that can generate the masking signal and time-multiplex on-chip, however this is still

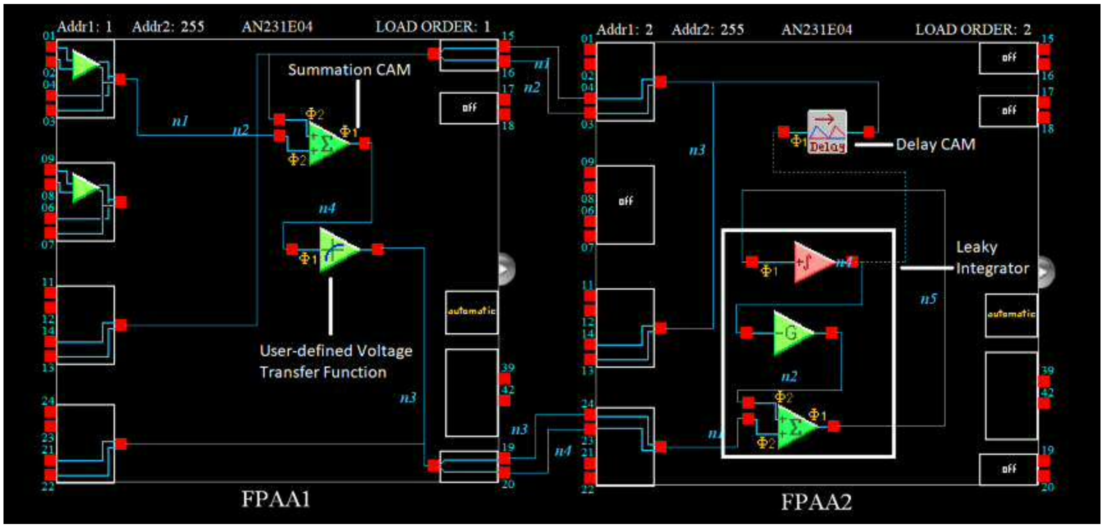

Fig. 1. Implementation of a Delay-Feedback reservoir using the Anadigm Designer 2 Tool. Here are two interconnected FPAA chips, labelled FPAA1 and FPAA2. The first chip contains a user-defined voltage transfer function, allowing for an arbitrary function to be implemented, and a summation CAM that sums the input and delayed feedback. The second chip contains an integrator with a leaky decay rate that defines the connectivity, and a delay CAM that adds a user defined delay between the input and output.

under investigation. At the moment, the input is pre-processed externally and loaded onto the FPAA to evaluate feasibility as a reservoir.

A powerful feature of the FPAA is its dynamic reconfigurability. This allows for all of the CAM parameters to be reconfigured on-line, allowing for a fully adaptable system. This could allow new RC architectures that change the RC's configuration and dynamics in real-time in order to tackle a wider range of computational problems.

IV. Conclusion

Here we outline the initial concept and prospects of a FPAA-based reservoir computing system. Testing the extended feasibility of the FPAA device is still underway and task benchmarking is still needed. The next steps are to evaluate the system on time series prediction data, such as the Nonlinear Autoregressive Moving Average (NARMA) tasks to undertake direct comparisons with other physical and simulated delay-feedback reservoirs.

A downside to current FPAA technology is that only a limited number of CABs can be fabricated within a single chip (typically four). This means the implementation of large interconnected networks on a single chip is somewhat unfeasible for now, and why the delay-feedback reservoir is chosen. However, it is possible to program chips individually and cascade them together to create increasingly complex RC systems.

References

[1] R. Pascanu, T. Mikolov, and Y. Bengio, "On the difficulty of training recurrent neural networks," in *International conference on machine learning*. PMLR, 2013, pp. 1310–1318.

[2] H. Jaeger, *Short term memory in echo state networks*. GMD-Forschungszentrum Informationstechnik, 2001.

[3] B. Schrauwen, D. Verstraeten, and J. Van Campenhout, "An overview of reservoir computing: theory, applications and implementations," in *Proceedings of the 15th european symposium on artificial neural networks. p. 471-482 2007*, 2007, pp. 471–482.

[4] G. Tanaka, T. Yamane, J. B. Héroux, R. Nakane, N. Kanazawa, S. Takeda, H. Numata, D. Nakano, and A. Hirose, "Recent advances in physical reservoir computing: A review," *Neural Networks*, 2019.

[5] L. Appeltant, M. C. Soriano, G. Van der Sande, J. Danckaert, S. Massar, J. Dambre, B. Schrauwen, C. R. Mirasso, and I. Fischer, "Information processing using a single dynamical node as complex system," *Nature communications*, vol. 2, no. 1, pp. 1–6, 2011.

[6] Y. Paquot, F. Duport, A. Smerieri, J. Dambre, B. Schrauwen, M. Haelterman, and S. Massar, "Optoelectronic reservoir computing," *Scientific reports*, vol. 2, no. 1, pp. 1–6, 2012.

[7] Anadigm. (2016) Anadigm quadapex development board. [Online]. Available: "https://anadigm.com/_doc/UM231004-K002.pdf"

Sculpting the Spin Ice for Computation

Arthur Penty and Gunnar Tufte
arthur.penty@ntnu.no

Norwegian University of Science and Technology (NTNU)

Ensembles of interacting nanomagnets known as Artificial Spin Ice (ASI) have become a promising new substrate for computation. Properties such as emergence and non-linear local interactions make it of particular interest for unconventional and material computation. Previously, we have proposed a method to represent and grow new ASI geometries, suited for use in an Evolutionary Algorithm (EA). Here we use our representation and evolution to further investigate towards computational properties including memory and classification. The richness of geometries found with sought computational properties indicates that ASI geometry is a fruitful tuning parameter for computational ASI systems.

1 Artificial Spin Ice

ASI [5] consist of a collection of nanomagnets, arranged such that they interact locally. The nanomagnets used are bi-stable, meaning we can represent the state of each magnet in an ASI as a 0 or 1, and thus the state of an ASI is the ensemble of these binary values. Though local interactions between the magnets are governed by a simple yet nonlinear formula, complex patterns form at higher levels. Most commonly, all the magnets are of the same, thus the different classes of patterns we can observe in ASI arise from different arrangements of the magnets. We refer to the arrangement of the nanomagnets as the ASI geometry.

A common approach to computing with ASI is to take some simple well-studied ASI geometry, excite or perturb it with an external magnetic field and observe the changes in its state. Work in this area focuses mostly on manipulating the strength and frequency of the external field, in order to achieve different computational properties. We take alternative approach, and instead modify the geometry of the ASI itself to achieve our computational goals.

2 Evolving geometry

In our previous work [3], we described a new representation for ASI geometries alongside a methodology to evolve these geometries, and demonstrated their performance on some simple problems via the use of an EA. For brevity, we do not give a detailed description of the representation here, only state that it consists of a small number of 'tile-like' building blocks which, through an iterative process, are used to build ASI geometries. Examples of this mapping can be seen in Fig. 1a and a complete description can be found in the aforementioned work. Once a geometry has been constructed, it can be passed to the flatspin ASI simulator [2] to evaluate its behaviour.

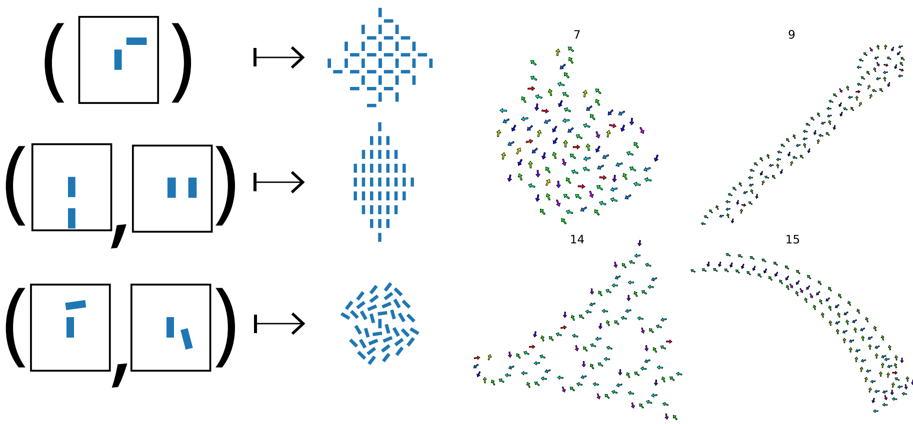

(a) ASI Mapping (b) Discovered Geometries

Fig. 1: (a) Examples of ASI building blocks or ‘tiles’, alongside the ASI geometry they produce under some configuration of the process described in [3].
(b) The final state of a small selection of the geometries found by novelty search, under after the input "0000". The number above each geometry indicates the number of distinct final states the 16 inputs can produce.

Our goal in this work is to use evolution to explore the richness of ASI dynamics in response to a global field input. In a similar vein to Jensen et al. [1], we observe how many different final states are produced when subjecting the ASI to different input series. In contrast to Jensen et al., we encode input as the angle of the sinusoidal global fields we apply. Inputs of 0 or 1 are mapped to a field applied at 0°or 90°respectively. Here we consider input bit strings of length four (which maps to four consecutive field application). For a given geometry, we independently perform each of the 16 possible 4-bit bit strings and observe the final state of the ASI. We can view the number of distinct final states a geometry can achieve, as a property of the intrinsic computation of the ASI acting on the input. If all inputs are mapped to one of two possible final states, this can be viewed as a 2-bin classifier. Conversely, if each input string leads to a unique final state then the ASI exhibits perfect memory on inputs of this length.

Employing our ASI geometry representation to generate new geometries, and the flatspin simulator [2] to evaluate them, we perform a novelty search to see how many geometries we can find that have different numbers of distinct end states. Fig. 1b shows some of the geometries found in this search. In fact, we were able to find a geometry for every possible number of distinct end states, i.e., for any $n \in [1, 16]$ we can supply a geometry that partitions all possible 4-bit inputs into n bins. We feel this is a very promising starting point for this route of tuning ASI geometries for specific kinds of computation, and we hope to further extend this to larger inputs soon.

Acknowledgements: This work was funded in part by the Norwegian Research Council IKTPLUSS project SOCRATES (Grant no. 270961), and in part by the EU FET-Open RIA project SpinENGINE (Grant no. 861618). Simulations were executed on the NTNU EPIC compute cluster[4].

References

1. Jensen, J.H., Folven, E., Tufte, G.: Computation in artificial spin ice. The 2018 Conference on Artificial Life: A Hybrid of the European Conference on Artificial Life (ECAL) and the International Conference on the Synthesis and Simulation of Living Systems (ALIFE) (30), 15–22 (2018). https://doi.org/10.1162/isal_a_00011, https://www.mitpressjournals.org/doi/abs/10.1162/isal_a_00011
2. Jensen, J.H., Strømberg, A., Lykkebø, O.R., Penty, A., Själander, M., Folven, E., Tufte, G.: Flatspin: A Large-Scale Artificial Spin Ice Simulator. arXiv:2002.11401 [cond-mat, physics:physics] (Feb 2020)
3. Penty, A., Tufte, G.: A Representation of Artificial Spin Ice for Evolutionary Search. ALIFE 2021: The 2021 Conference on Artificial Life (07 2021). https://doi.org/10.1162/isal_a_00436, https://doi.org/10.1162/isal_a_00436, 99
4. Själander, M., Jahre, M., Tufte, G., Reissmann, N.: EPIC: An Energy-Efficient, High-Performance GPGPU Computing Research Infrastructure. arXiv:1912.05848 [cs] (Dec 2019)
5. Skjærvø, S.H., Marrows, C.H., Stamps, R.L., Heyderman, L.J.: Advances in artificial spin ice. Nature Reviews Physics **2**(1), 13–28 (Jan 2020). https://doi.org/10.1038/s42254-019-0118-3, http://www.nature.com/articles/s42254-019-0118-3