Computation in artificial spin ice

Johannes H. Jensen[1], Erik Folven[2] and Gunnar Tufte[1]

[1]Department of Computer Science,
[2]Department of Electronics and Telecommunications,
Norwegian University of Science and Technology, Trondheim, Norway
johannes.jensen@ntnu.no

Abstract

We explore artificial spin ice (ASI) as a substrate for *material computation*. ASI consists of large numbers of nanomagnets arranged in a 2D lattice. Local interactions between the magnets gives rise to a range of complex collective behavior. The ferromagnets form large networks of nonlinear nodes, which in many ways resemble artificial neural networks. In this work, we investigate key computational properties of ASI through micromagnetic simulations. Our nanomagnetic system exhibits a large number of reachable stable states and a wide range of available dynamics when perturbed by an external magnetic field. Furthermore, we find that the system is able to store and process temporal input patterns. The emergent behavior is highly tunable by varying the parameters of the external field. Our findings highlight ASI as a very promising substrate for *in-materio* computation at the nanoscale.

Introduction

Intelligent systems in nature are closely coupled to physics. Through bottom-up exploration and exploitation of physical processes, evolution has found ways to achieve self-organized computation. Such natural computation ultimately results in intelligent behavior (Mainzer, 2007). Furthermore, natural computation is extremely efficient: Our brain contains billions of processing elements (neurons) but consumes only 20W.

Artificial intelligent systems, e.g. artificial neural networks, are abstracted far away from the physical and chemical domain. As such, physical properties are not exploitable for computation. Consequently, running these abstract models is inherently inefficient and requires massive server farms consuming megawatts of power.

Material computation places physics back in the front seat and views computation as an inherently physical process. Within a physical system (the material) lies an inherent capacity for computation. The goal is to find ways to exploit this inherent computing power. If accomplished, we can build extremely efficient computing devices based on principles found in biological systems: vast parallelism, self-organization, robustness and adaptation (Stepney, 2008).

The key principles of material computation date back to the early days of cybernetics (Ashby, 1960; Pask, 1959) and artificial life (Langton, 1990): a large number of nodes, nonlinearity, local interactions and rich dynamics. More recently the field of reservoir computing (Jaeger, 2001) has proven that complex physical systems with such properties can be readily exploited, by training a readout layer to map system output to the target problem (Dale et al., 2017; Jensen and Tufte, 2017; Sillin et al., 2013).

Hopfield (1982) was early to establish links between neural network models and physics. Hopfield networks are defined in terms of an energy function whose many local minima represent memorized states. Recently there has been an increased interest in energy-based neural network models since they are physically plausible learning architectures (Scellier and Bengio, 2017).

Hopfield energy provides a direct link between neuronal computation and ferromagnetic systems. The energy function is isomorphic with an Ising model where atomic spins take the place of neurons, while exchange coupling between spins is analogous to synaptic connections. Hopfield showed that important properties for computation spontaneously arise in such systems with many nonlinear nodes.

Artificial spin ice (ASI) is a class of ferromagnetic metamaterial which consists of large arrays of coupled nanomagnets. ASI exhibits key properties for material computation: a large number of nonlinear elements whose local interaction gives rise to complex collective behavior. Furthermore, the physical fabrication of ASI systems has been well-established over the last decade. It is thus an intriguing substrate for neuronal material computation.

In this work we investigate key properties for material computation in an ASI substrate. Through detailed micromagnetic simulations, we explore the range of dynamics available and the number of reachable states when ASI is subject to an oscillating magnetic field. Finally we demonstrate how the system can be exploited for temporal pattern classification.

This paper is organized as follows: we begin with a short introduction to the field of artificial spin ice. Next, we dis-
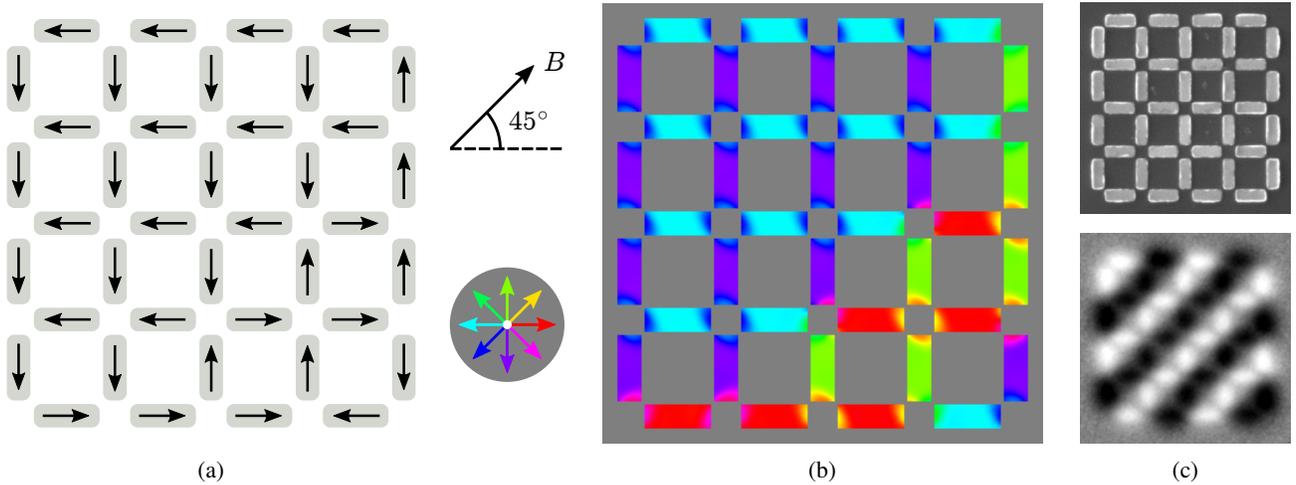
Figure 1: Square artificial spin ice consists of horizontal and vertical nanomagnets arranged in a 'square' pattern. For this work we study the 4x4 square spin ice which consists of 40 nanomagnets (a). Since the nanomagnets are single-domain, the internal magnetization will align left/right for horizontal magnets and up/down for vertical magnets (b). Colors indicate the direction of magnetization as shown in the color map. We perturb the array with an external magnetic field B at an angle of $45°$ as shown. Microscopic imaging of our fabricated spin ice (c) enables a view of the geometry (top) and magnetization (bottom) of each individual nanomagnet. The size of each magnet is 220x80 nm.

cuss how ASI relates to established models of computation. We introduce our methods and simulation experiments, followed by results and discussion. Finally we conclude our findings and discuss future work.

Artificial spin ice

Artificial spin ice is a class of metamaterial which consists of nanoscale ferromagnetic islands (nanomagnets) arranged in a 2D lattice. Figure 1a shows the spin ice system used in this study, which consists of 40 nanomagnets arranged in a 4x4 square lattice.

For nanomagnets smaller than some critical dimension, the individual spin moments will tend to align in the same direction, i.e. a *single-domain* state. When the magnets are elongated, the ground state magnetization will be oriented homogeneously in either of the two directions defined by the long-axis of the magnet. Figure 1b shows how the internal magnetization of each nanomagnet is oriented in a single-domain ground state. Magnetization reversal, or *switching* in such nanoscale magnets is an inherently nonlinear process (Stoner and Wohlfarth, 1948).

Nanomagnets in the lattice will interact through local dipolar coupling. The interactions depend on the geometrical arrangement of the lattice. This makes it possible to create a geometrically frustrated system, i.e. a system in which all competing interactions cannot be satisfied at the same time. Many geometric patterns giving rise to geometrical frustration have been demonstrated. In this work we focus on the 'square' lattice as shown in Figure 1.

Originally, the field of ASI was developed to build nan-

otechnological model systems for fundamental studies of geometrical frustration (Wang et al., 2006). It has since grown to encompass a wide variety of phenomena ranging from effects of quenched disorder, thermally activated dynamics, microwave frequency responses, magnetotransport properties, and magnetic phase transitions (Marrows, 2016).

Owing to the semiconductor industry, the fabrication routes for artificial spin ice devices are already highly developed. With electron beam lithography, extended arrays consisting of millions of elements can readily be fabricated. The variety of geometrical arrangements that can be realized is only limited by the creativity of the designer and a plethora of different lattices have been created and explored in existing literature (Heyderman and Stamps, 2013).

The development of microscopy techniques with lateral resolution at the nanoscale and magnetic contrast enables direct imaging of the magnetization of individual nanomagnet elements (Figure 1c). At the macroscale, the collective state of the array can be observed using well-established magnetometry techniques. Reading techniques beyond the lab include magnetic tunnel junction based approaches similar to those found in Magnetic RAM and conventional hard drives.

There are many ways to manipulate the nanomagnet elements within an ASI. External magnetic fields is a well-established approach, applied either globally to the entire array or locally to specific areas, e.g. via current-carrying electrical wires. Other possibilities include current-induced torques (Brataas et al., 2012), optically induced switching (Le Guyader et al., 2015) and using a scanning probe to manipulate individual nanomagnets (Gartside et al., 2018).

Recent developments of GPU accelerated simulation frameworks combined with increased computational resources have opened for simulation studies of extended nanomagnet ensembles, which greatly aids the exploration of novel artificial spin ice configurations.

Computation in artificial spin ice

ASI exhibits key properties for material computation, i.e. a large number of nodes, nonlinearity and local interactions. However, the question of how ASI relates to models of computation needs to be answered. Here we discuss the relation between ASI and one model of computation: artificial neural networks (ANNs).

ASI has a range of properties which are analogous to ANNs. Like neural networks, ASI consists of a large number of nonlinear nodes (nanomagnets) which are connected together in a network (dipolar coupling). Magnets are non-volatile devices, so each node exhibits long term memory. As with neural networks, computation in ASI is thus closely coupled with memory.

Network dynamics is the result of local interactions between neighboring magnets, much like recurrent neural networks. Geometry imposes certain limitations on the network topology, e.g. fully connected networks are not realizable in ASI. In this aspect ASI shares similarities with cellular neural networks (Chua and Yang, 1988).

A major difference between ASI and ANNs is the absence of intrinsic synaptic plasticity: the coupling between magnets is determined by their geometrical arrangement which is fixed. ASI is thus analogous to a network with fixed weights. The coupling strength can however be modified externally, e.g. through local magnetic fields. Reservoir computing has proven that even random fixed-weight networks can be exploited to solve many useful problems (Jaeger, 2001).

Microstates and macrostates

The state of a physical system can be understood in terms of hierarchies, i.e. the level of observation. We can distinguish between the *microstate* and the *macrostate*. The microstate is the state at the most detailed level, i.e. down to the smallest particle. Usually we do not have access to the microstate. Typically we only have access to the macrostate: the state of the system observed at some higher level. As observers, we are free to choose how and at what level we measure the macrostate.

In the case of ASI, the microstate is defined by the magnetic moments of the atoms within a nanomagnet. In simulation we can directly observe the microstate, something which is not possible in the real physical system. Figure 1b shows the microstate of our simulated 4x4 square spin ice.

Regarding the macrostate of ASI, a natural choice is to define a macrostate based on the average magnetization within a nanomagnet. For the single domain magnets used here,

the magnetization will align in one of two directions, i.e. a *binary* macrostate. The macrostate of the entire array can thus be represented with N bits, where N is the number of nanomagnets. Figure 1a shows the macrostate derived from the corresponding microstate in Figure 1b. The macrostate as can be observed in the real physical system can be seen in the bottom of Figure 1c.

Macrostates are in general *degenerate*, i.e. there are many microstates which map to the same macrostate. Information about the true state of the system is therefore hidden. States that look identical at the macro level may in reality be different at the micro level.

Energy-based computation

Energy-based models such as Hopfield networks define computation as movement through an energy landscape. Each valley is a local energy minimum and represent stable ground states of the system. Given an arbitrary initial state, the system will settle into the nearest valley. If there are many such valleys, the system has a large number of stable states.

If we perturb the system with sufficient force (input), the system may escape the valley and transition to a nearby state. This movement through state space is a form of intrinsic physical computation. Computation is thus closely linked to dynamics, i.e. complex dynamics results in complex computations (Langton, 1990).

When ASI is subject to a cyclic external magnetic field, complex switching behavior can occur. The behavior will depend critically on the strength of the magnetic field. If the field is too weak, there is not enough energy to leave the local energy minimum so no switching occurs. If the field is too strong, all the magnets will simply follow the field and we get trivial switching behavior. However, when the field strength is at some critical value, the local magnetic fields around each magnet will affect the switching of neighboring magnets. At this critical field strength we may find complex dynamics. Gilbert et al. (2015) demonstrated experimentally that for slowly varying fields whose strength is just above the array coercivity, an ASI will go through several transient states before settling in a stable ground state.

Little is however known about the switching dynamics of ASI at high frequencies. It seems likely that increasingly complex dynamics occur at higher frequencies where phenomena such as spin waves come into play. Indeed, studies of magnetic resonance has revealed complex frequency dependence of spin ice systems in the GHz range (Jungfleisch et al., 2016, 2017). Furthermore, a system of two interacting dipoles exhibits chaotic dynamics when subject to a time-dependent external magnetic field (Urzagasti et al., 2015).

Methods

We have argued that artificial spin ice exhibits key properties for material computation. Several properties are fulfilled by
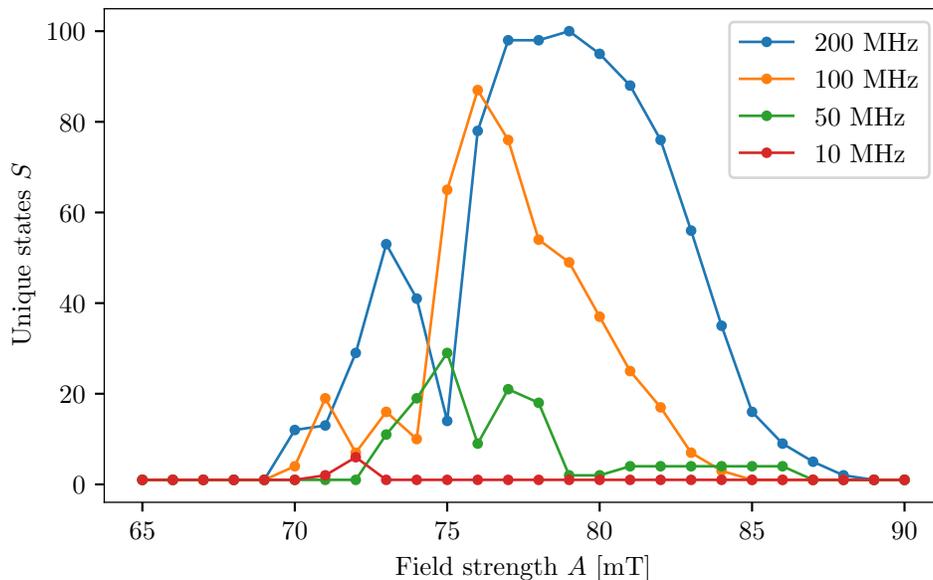
Figure 2: Number of unique states S visited as a function of the external field strength A, for four different frequencies.

definition: a large number of nodes, nonlinearity and local interactions. However, the availability of complex dynamics within ASI is still largely unexplored. We need a reliable way to excite and control such complex dynamics.

Here we investigate magnetization dynamics when ASI is perturbed by an external magnetic field of high frequency. Next we demonstrate how such complex dynamics can be exploited for computation.

As subject of study we use the 4x4 square spin ice depicted in Figure 1 consisting of 40 permalloy nanomagnets. The size and spacing of the magnets is similar to previous studies (Wang et al., 2006): each magnet is 220x80 nm with a thickness of 25 nm, and the lattice spacing between each magnet is 320 nm.

We start with an initially polarized array such that all horizontal magnets point to the right and vertical magnets point up. The initial state is easily obtained by first saturating the array with a strong magnetic field at $45°$ which is then reduced to zero.

We perturb the array with a time-varying external magnetic field $B(t) = A \sin(\omega t)$ with amplitude A and frequency ω. The field is applied at a constant angle of $45°$ as shown in Figure 1.

The magnets in our experiments are single domain, so we adopt a binary macrostate based on the average magnetization. For horizontal magnets, let the state be 1 if the magnetization points to the right and 0 if it points to the left. Similarly for vertical magnets, the state is 1 if the magnetization points up and 0 if it points down. Since the array consists of 40 magnets, the macrostate of the entire array can be represented as a 40 bit vector. Hence the system has a state space of 2^{40} unique states.

All experiments were conducted using the MuMax3 micromagnetic simulator (Vansteenkiste et al., 2014). Key material parameters used are $M_{sat} = 860 \times 10^3$ A/m, $A_{ex} = 13 \times 10^{-12}$ J/m and $\alpha = 0.01$. A lateral cell size of 5x5 nm^2 was used throughout, which is less than the exchange length ($L_{ex} = 5.3$ nm).

Experiment 1: Complex dynamics

First we investigate the dynamics of the array when the sinusoidal external field has constant amplitude and frequency. In the following, we investigate the impact of field amplitude A and frequency ω on the dynamics of the system.

We perturb the array with 100 periods of the external field and sample the state of the system at the end of each period. We then count the number of unique states S visited during this time ($1 \leq S \leq 100$). A large number of unique states indicates complex dynamics while a low number indicates stability. A hallmark of chaos is aperiodic long-term behavior where state space trajectories never repeat, hence a large number of unique states is an indication of chaotic dynamics.

For weak fields, we expect $S = 1$ since none of the magnets will switch. For strong fields, we also expect $S = 1$ as all magnets will switch in unison with the field, and return back to the same polarized state after one period. Only for intermediate field strengths close to the array coercivity H_c can we hope to find complex switching dynamics and consequently $S \gg 1$. For our system we estimate $H_c \approx 75$ mT measured at $45°$.

We vary the field amplitude A around H_c and count the number of states for each A. This sweep is repeated for a set of frequencies.
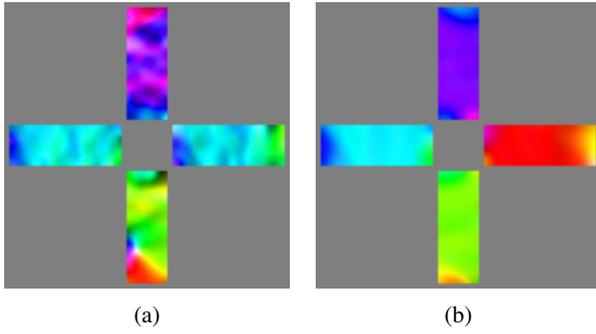
(a)   (b)

Figure 3: Snapshots of the microstates for 200 MHz (a) and 100 MHz (b). The color map is identical to Figure 1.

Results and discussion

Figure 2 shows the number of unique states S as a function of the external field strength A. We plot the number of states for four different frequencies: 10 MHz, 50 MHz, 100 MHz and 200 MHz.

At 10 MHz we see at most 6 unique states which is in agreement with Gilbert et al. (2015). When the frequency is increased to 50 MHz, the number of states reaches a maximum of 29 for $A = 75$ mT. At 100 MHz there is a significant increase in number of states with a maximum of 87 for $A = 76$ mT. As we increase frequency further to 200 MHz, we observe a saturation in the number of states ($S = 100$) for $A = 79$ mT.

As expected, the bell-like curves are all centered around the array coercivity $H_c \approx 75$ mT. A general trend is that the number of unique states increases with frequency. The curves also become wider with increased frequency.

The large number of unique states seen at frequencies 100 MHz and above is the product of long transients. It is thus likely that dynamics are chaotic in these cases. The results also indicate that the dynamics get more chaotic as frequency is increased. Videos of the dynamics are available on our website[1].

These results demonstrate that ASI has a large number of *reachable* macrostates. Furthermore, the states can be reached by the straightforward application of a cyclic magnetic field, as long as the field strength is close to a critical value and the frequency is sufficiently high.

We get a better understanding of these results by examining the microstate. Figure 3 shows snapshots of the microstate taken at the end of the field cycle, i.e. at zero amplitude. Figure 3a shows the microstate at 200 MHz, where we can see that spin wave dynamics have not yet settled. These micro level dynamics become very turbulent if frequency is increased further to 1 GHz (not pictured). Hence at high frequencies, the spins within the nanomagnet have not yet aligned in one direction. In other words, the system

_____
[1]https://www.ntnu.edu/socrates/magnets
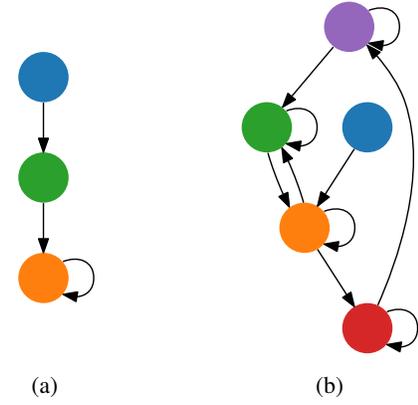

(a)   (b)

Figure 4: Example state transition graphs for 10 MHz (a) and 50 MHz (b).

has not reached a stable equilibrium and consequently the macrostate will be in flux. Figure 3b shows the microstate at 100 MHz, where we see that the spin waves have mostly settled at the end of the field cycle, resulting in a stable macrostate.

What effect do micro level dynamics such as spin waves have on dynamics at the macro level? An interesting question is whether macro level dynamics are memoryless, i.e. does the next macrostate depend only on the current macrostate? We answer this with an analysis of the state transition graphs derived from the observed macro level dynamics (Figure 4). Here nodes represent macrostates and edges are the transitions between them.

At low frequencies (10 MHz) we find that all nodes have at most one outgoing edge, i.e. the dynamics are memoryless. An example is shown in Figure 4a where the system settles in a stable state after a transition through an intermediate state.

At higher frequencies (50 MHz and above) most graphs contain nodes with more than one outgoing edge, as shown in Figure 4b. Here the next macrostate can not be determined without knowing the history of previous macrostates. This is information which is hidden in the microstate.

This suggests that the microstate provides a means of information storage, which is indirectly observed in the history dependent dynamics at the macro level. Indeed it has been shown that chaos both generates and stores information (James et al., 2014).

In summary, these results demonstrate that rich micro level dynamics are available in the ASI system. Dynamics at the micro level give rise to phenomena at the macro level, i.e. a large number of distinct macrostates and complex state transition patterns. Crucially, the dynamics appear to be highly tunable, e.g. by varying the amplitude and/or frequency of an external magnetic field.
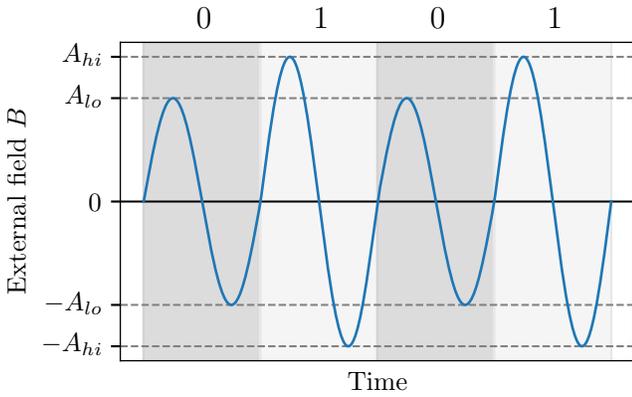
Figure 5: Bit strings are encoded into the external magnetic field with amplitude modulation. Each bit corresponds to one cycle of the external field, where 1 maps to an amplitude A_{hi} while 0 maps to an amplitude A_{lo}. The figure shows how the bit string 0101 is encoded.

Experiment 2: Temporal pattern classification

Results from the first experiment show that information about state history is embedded in the microstate, which suggests that the system may be able to store and process temporal input.

Let us assume that transitions between macrostates is sensitive not only to the history of past states, but also to the history of the external field. We can exploit this property for computation by encoding input as part of the time-varying external field.

For each input, does the system end up in a unique macrostate? If so, the system essentially acts as *memory*, transforming a temporal input to a unique spatial state. The function is that of maximum discrimination: each input is mapped to a unique output state. Hence, information in the input is preserved by the system.

On the other hand, if the system ends up in only a handful of states, the functionality is some form of *classification*. The input-output relations are many to one, i.e. many inputs map to the same output state. This mapping may be arbitrarily complex, but information about the input is always lost in the mapping process.

What about the case in-between these two extremes? Here the functionality is a mixture of memory and classification where some information is retained and some information is lost in the mapping.

Can we find these different modes of computation in square artificial spin ice? To test this, we consider as input bit strings of length $N = 1$ to $N = 8$. For each N we apply all possible 2^N inputs and record the final macrostate of the array. We then count the number of unique such states to determine what mode of computation is performed: memory, classification or a mixture of the two.

To encode the bit strings we employ amplitude modula-

tion of the external magnetic field as shown in Figure 5. Each bit corresponds to one cycle of the external field, where 1 maps to an amplitude A_{hi} while 0 maps to an amplitude A_{lo}.

We set the frequency of the external field to 100 MHz in order to obtain both complex dynamics and a stable macrostate at the end of each period. For the current study we fix $A_{lo} = 70$ mT and vary A_{hi} in the range where complex switching dynamics was found in the previous experiment, i.e. 70 mT $< A_{hi} \leq 84$ mT.

Results and discussion

Figure 6 shows the number of unique final states S as a function of the number of bits N in the input string. For comparison we plot the function 2^N which is the number of different input values and thus the theoretical maximum. We plot the number of unique states for a selection of A_{hi} values which resulted in distinct behavior, namely 76 mT, 79 mT, 81 mT and 84 mT.

A general trend for all the values of A_{hi} is that the number of states increases with the number of input bits. The rate of increase is however quite different for distinct values of A_{hi}. Another observation is that the number of states do not appear to saturate.

Recall from the first experiment that $A = 76$ mT was the amplitude which produced the highest number of states at 100 MHz. Here we see that this particular value for A_{hi} also results in the highest number of unique final states, closely tracking the theoretical maximum 2^N. In this regime the system acts primarily as memory, mapping almost all temporal input patterns to unique macrostates. At the maximum $N = 8$ the system maps the 256 different input patterns to 226 different states.

On the other extreme we have $A_{hi} = 84$ mT which results in only a handful of unique states. In this regime the system acts as a form of classifier. When the number of input bits is 8, the system maps the 256 different input patterns to 13 distinct states.

Interestingly, we find that certain values of A_{hi} result in functionality somewhere in between memory and classification. For $A_{hi} = 79$ mT we can see that the number of states follows roughly in the middle of the previous extremes, while for $A_{hi} = 81$ mT the state count is somewhat lower.

We may understand these modes of intrinsic computation in terms of dynamical systems theory. Classification is equivalent to entering an attractor, where the number of attractors is equal to the number of classes. Memory can be explained by chaos where sensitivity to initial conditions means that every distinct input will result in a unique trajectory through state space (Crook, 2007; Shaw, 1981).

The sensitivity of the system will determine how quickly nearby trajectories diverge and consequently whether similar input patterns will end up in the same macrostate after all
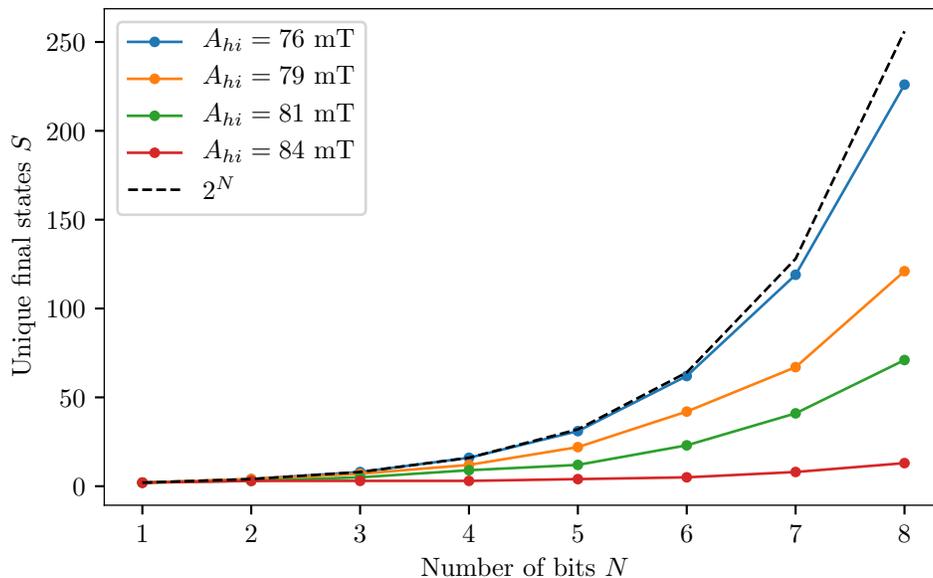
Figure 6: Number of unique final states S as a function of number of bits N in the input string (100 MHz input frequency).

the bits have been applied. By adjusting the field amplitude A_{hi} we are essentially tuning the sensitivity of the system, resulting in different modes of computation.

These results demonstrate only one way to exploit ASI for computation with a temporal input encoding. Here we are exploiting the rich repertoire of dynamics available in the system. Combined with the large number of stable states, the system is able to effectively distinguish between temporal patterns. Furthermore, we show that the sensitivity of discrimination is highly tunable which yields different modes of computation.

Conclusion and future work

In this work we have explored artificial spin ice (ASI) as a highly promising substrate for material computation. The ferromagnetic material exhibits key intrinsic properties for computation: a large number of nodes, nonlinearity, local interactions and rich dynamics. Fabrication methods for ASI are already highly developed, making such computational properties readily available *in-materio*.

Through micromagnetic simulations of square ASI, we have found a large number of reachable stable states and a wide range of available dynamics when ASI is perturbed by an oscillating magnetic field. Furthermore, the dynamics are highly tunable by varying the amplitude and frequency of the magnetic field.

We have also demonstrated how the complex dynamics and many stable states in ASI can be exploited for temporal pattern classification. By tuning the sensitivity of the system, different modes of computation can be obtained: memory, classification or a mixture of the two.

The results in this work have been derived from micro-

magnetic simulations. A natural next step is to replicate these results experimentally with physical realizations of the spin ice. As shown in Figure 1c, the physical parameters of the ASI are all well within limits of what can be fabricated and measured in our lab today.

Our experiments so far have been purely deterministic without noise. A key question is how the system behaves in the presence of noise, e.g. thermal noise at room temperature or electrical noise from measurement equipment. Results indicate that the system may be very sensitive to small changes, i.e. chaotic dynamics will be susceptible to noise. Hence, it is critical to obtain a measure of this sensitivity.

The ASI system used in this study has been rather small, consisting of only 40 magnets. However, scaling up to millions of magnets can easily be achieved in physical realizations. A key question is then how scale affects the behavior of larger arrays. Micromagnetic simulations of such large systems is not computationally feasible, hence it is crucial to establish simulation models at higher levels of abstraction.

The focus of this work has been on the basic computational properties of ASI. As such, we have not focused on any specific application in our experiments. Future work will include the application of techniques such as reservoir computing to exploit ASI for useful tasks. Reservoir computing in spin ice could enable robust, massively parallel magnetic processing at the nanoscale.

The work presented herein only scratches the surface of what is possible with ASI systems. There is a wealth of parameters worth exploring, e.g. magnet size, shape, spacing and geometry, together with methods for perturbation and observation of dynamics. Furthermore, developing ways to efficiently exploit the intrinsic computing power in such sys-

tems is critically important. The field of ASI computation is ripe for exploration, towards vastly parallel and energy efficient computing substrates.

Acknowledgements

References

Ashby, W. R. (1960). *Design for a Brain: The Origin of Adaptive Behavior*. Chapman & Hall, London, England.

Brataas, A., Kent, A. D., and Ohno, H. (2012). Current-induced torques in magnetic materials. *Nature Materials*, 11(5):372–381.

Chua, L. and Yang, L. (1988). Cellular neural networks: theory. *IEEE Transactions on Circuits and Systems*, 35(10):1257–1272.

Crook, N. (2007). Nonlinear transient computation. *Neurocomputing*, 70(7-9):1167–1176.

Dale, M., Miller, J. F., and Stepney, S. (2017). Reservoir Computing as a Model for In-Materio Computing. In Adamatzky, A., editor, *Advances in Unconventional Computing. Emergence, Complexity and Computation, vol 22.*, chapter 22, pages 533–571. Springer, Cham.

Gartside, J. C., Arroo, D. M., Burn, D. M., Bemmer, V. L., Moskalenko, A., Cohen, L. F., and Branford, W. R. (2018). Realization of ground state in artificial kagome spin ice via topological defect-driven magnetic writing. *Nature Nanotechnology*, 13(1):53–58.

Gilbert, I., Chern, G.-W., Fore, B., Lao, Y., Zhang, S., Nisoli, C., and Schiffer, P. (2015). Direct visualization of memory effects in artificial spin ice. *Physical Review B*, 92(10):104417.

Heyderman, L. J. and Stamps, R. L. (2013). Artificial ferroic systems: novel functionality from structure, interactions and dynamics. *Journal of Physics: Condensed Matter*, 25(36):363201.

Hopfield, J. J. (1982). Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences*, 79(8):2554–2558.

Jaeger, H. (2001). The "echo state" approach to analysing and training recurrent neural networks - with an Erratum note. Technical Report 148, German National Research Center for Information Technology.

James, R. G., Burke, K., and Crutchfield, J. P. (2014). Chaos forgets and remembers: Measuring information creation, destruction, and storage. *Physics Letters, Section A: General, Atomic and Solid State Physics*, 378(30-31):2124–2127.

Jensen, J. H. and Tufte, G. (2017). Reservoir computing with a chaotic circuit. In *Proceedings of the 14th European Conference on Artificial Life ECAL 2017*, pages 222–229, Cambridge, MA. MIT Press.

Jungfleisch, M. B., Sklenar, J., Ding, J., Park, J., Pearson, J. E., Novosad, V., Schiffer, P., and Hoffmann, A. (2017). High-Frequency Dynamics Modulated by Collective Magnetization Reversal in Artificial Spin Ice. *Physical Review Applied*, 8(6):064026.

Jungfleisch, M. B., Zhang, W., Iacocca, E., Sklenar, J., Ding, J., Jiang, W., Zhang, S., Pearson, J. E., Novosad, V., Ketterson, J. B., Heinonen, O., and Hoffmann, A. (2016). Dynamic response of an artificial square spin ice. *Physical Review B*, 93(10):100401.

Langton, C. G. (1990). Computation at the edge of chaos: Phase transitions and emergent computation. *Physica D: Nonlinear Phenomena*, 42(1-3):12–37.

Le Guyader, L., Savoini, M., El Moussaoui, S., Buzzi, M., Tsukamoto, A., Itoh, A., Kirilyuk, A., Rasing, T., Kimel, A. V., and Nolting, F. (2015). Nanoscale sub-100 picosecond all-optical magnetization switching in GdFeCo microstructures. *Nature Communications*, 6(1):5839.

Mainzer, K. (2007). *Thinking in Complexity*. Springer Berlin Heidelberg, Berlin, Heidelberg, 5 edition.

Marrows, C. (2016). Experimental Studies of Artificial Spin Ice. *arXiv preprint arXiv:1611.00744*.

Pask, G. (1959). Physical Analogues to the Growth of a Concept. In *Mechanization of Thought Processes, Symposium*, pages 877–922.

Scellier, B. and Bengio, Y. (2017). Equilibrium Propagation: Bridging the Gap between Energy-Based Models and Backpropagation. *Frontiers in Computational Neuroscience*, 11:24.

Shaw, R. (1981). Strange Attractors, Chaotic Behavior, and Information Flow. *Zeitschrift für Naturforschung A*, 36(1):81–87.

Sillin, H. O., Aguilera, R., Shieh, H.-H., Avizienis, A. V., Aono, M., Stieg, A. Z., and Gimzewski, J. K. (2013). A theoretical and experimental study of neuromorphic atomic switch networks for reservoir computing. *Nanotechnology*, 24(38):384004.

Stepney, S. (2008). The neglected pillar of material computation. *Physica D: Nonlinear Phenomena*, 237(9):1157–1164.

Stoner, E. C. and Wohlfarth, E. P. (1948). A Mechanism of Magnetic Hysteresis in Heterogeneous Alloys. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 240(826):599–642.

Urzagasti, D., Becerra-Alonso, D., Pérez, L. M., Mancini, H. L., and Laroze, D. (2015). Hyper-chaotic Magnetisation Dynamics of Two Interacting Dipoles. *Journal of Low Temperature Physics*, 181(5-6):211–222.

Vansteenkiste, A., Leliaert, J., Dvornik, M., Helsen, M., Garcia-Sanchez, F., and Van Waeyenberge, B. (2014). The design and verification of MuMax3. *AIP Advances*, 4(10):107133.

Wang, R. F., Nisoli, C., Freitas, R. S., Li, J., McConville, W., Cooley, B. J., Lund, M. S., Samarth, N., Leighton, C., Crespi, V. H., and Schiffer, P. (2006). Artificial 'spin ice' in a geometrically frustrated lattice of nanoscale ferromagnetic islands. *Nature*, 439(7074):303–306.