Journal of Computational Science 2 (2011) 279–285

Contents lists available at ScienceDirect

Journal of Computational Science

journal homepage: www.elsevier.com/locate/jocs

Using a custom-FPGA architecture to extend the scale of atomistic magnetic spin simulations

David Cortie[a,b,*], John Pillans[b]

[a] Institute for Superconducting and Electronic Materials, University of Wollongong, Wollongong, New South Wales, Australia
[b] Australian Nuclear Science and Technology Organisation, New Illawara Road, Sydney, New South Wales, Australia

ARTICLE INFO

Article history:
Received 16 February 2011
Received in revised form 21 April 2011
Accepted 27 April 2011
Available online 19 May 2011

Keywords:
Field-programmable-gate-array
Magnetism
Modelling
Ising model

ABSTRACT

New computational solutions are required to understand how atomic-scale properties affect magnetic behaviour at micrometer dimensions. We describe a field-programmable-gate-array (FPGA) based simulation of a dilute antiferromagnet with a large number of Ising spins using Glauber dynamics. The simple atomic model qualitatively reproduces experimental findings when scaled up to sufficiently large spatial dimensions, and provides insight into the finite size thresholds separating nanoscale from microscale behaviour. A real-time visualisation module was used to study the dynamics of the fractal domain structure and non-exponential relaxation mechanism. A performance comparison with contemporary GPU and CPU implementations suggests that a FPGA route is a competitive alternative.

1. Introduction

Historically the study of magnetism in solids has motivated innovative approaches in mathematics, physics and computational science because simple experimental facts have complicated microscopic origins. For example, the origin of a permanent magnet's properties was only satisfactorily understood in 1944 when Onsager proved that atomic spins with Ising-behaviour could undergo a disorder–order phase transition into a stable ferromagnetic state [17]. To this day, significant gaps remain in contemporary scientific knowledge about magnetism, particularly regarding the properties of magnetically ordered materials that are grown into artificial geometries such as nano-wires or atomic-layer thin films. Computational modeling has proved to be a vital instrument for resolving many questions, but the diverse range of time scales (picoseconds to years) and spatial scales (nm to cm) [19] has necessitated multi-scale modeling which relies on a ladder of techniques [10]. For example, there is a fissure between discrete atomistic and continuous micromagnetic modeling which hinders the development of spintronic technology operating between these length scales. Large scale atomistic simulations are highly desirable, but the computational costs for systems that reach from atomic to micrometer scales are unattainable for first-principal calculations or classical Landau–Lifshitz–Gilbert (LLG) dynamics for discrete atomic spins. Slow creep dynamics of large spin systems have been studied using the Monte Carlo formalism derived from statistical mechanics, but the computational cost is prohibitive [15]. To increase the number of spins that can be simulated in a reasonable time frame, a series of computational approaches have been explored including custom-built computers [4], supercomputers [15] and, more recently, the use of graphical-process units (GPU's) [1,2,18]. Here we discuss an alternative solution using a field-programmable-gate-array (FPGA) implementation of the Ising Model with a Monte Carlo update algorithm.

2. FPGA implementations of scientific models

The field-programmable-gate-array is a compromise between the speed of an application specific integrated circuit (ASIC) and the re-programmable nature of software running on a CPU or GPU. Unlike an ASIC, the underlying FPGA architecture is hardware re-programmable allowing the reconfiguration of logic gates for specific computational tasks. Due to this unique capability, FPGA's are currently topical in the field of adaptive computation, cryptoanalysis and neural networks [20] but they have seldom been applied to scientific simulations of magnetism [6]. They are ideally suited to Ising Model Monte Carlo programs owing to their built-in parallelism and the binary nature of Ising spins. In this work we developed a single-chip FPGA architecture to study the Ising spin system dynamics on a simple cubic lattice using parallelised calculations. Large 3D lattice sizes could be simulated including $32 \times 32 \times 32$, $128 \times 128 \times 128$ and $512 \times 512 \times 128$ but increasing lattice size introduced constraints on the number of parallel processes which could be run. A simple system, parameterised in the

* Corresponding author. Tel.: +61 2 9717 3600.
E-mail address: dcr@ansto.gov.au (D. Cortie).

1877-7503/$ – see front matter.
doi:10.1016/j.jocs.2011.04.001

dimensions of the array, was developed using a low level hardware development language which allowed efficient utilisation of resources within the inexpensive Xilinx Spartan 3E 1200 FPGA. A real-time visualisation was included in the system without impact on computational performance, and several metrics were computed and streamed to an ordinary personal computer. To ensure numerical accuracy, the FPGA-simulation was first checked against Onsager's exact solution which analytically describes the simpler 2D Ising ferromagnet. We then applied the computational system to the case of a diluted antiferromagnet (DAFF) system which is a prototypical example of a random-field antiferromagnet. This material class displays a rich set of phenomena which cannot be solved analytically or described in a mean-field approach, but are technologically important for exchange bias layers in spin valves [14]. The FPGA architecture allowed the simulation to access larger spatial dimensions and time scales than have previously been studied.

3. Implementation

A Monte Carlo procedure with a heat-bath algorithm for determining spin flips was used to study the ensemble properties of an Ising spin system as a function of temperature and field. Initially the virtual spins are placed at regular points on a cubic lattice and given a random direction. To simulate dilution, certain spins are replaced by non-magnetic defect sites. For the FPGA model, we implement a bitwise representation so that 01 represents spin up, 00 represents spin down and 10 represents a defect site. The array of spin data is stored in a single long shift register, thereby mapping the 3D cubic lattice of spins into a 1D linear array. At each time step (Monte Carlo step), a sweep is made through the entire array of spins. During the sweep, the energy is calculated for each spin using the Hamiltonian energy equation for the dilute antiferromagnet in an external field [15]:

$$H = -J\sum_{ij}\varepsilon_i\varepsilon_j\sigma_i\sigma_j - \sum_i \mu B_z \varepsilon_i \sigma_i \tag{1}$$

A detailed explanation of the terms above is given in Appendix A. On an FPGA, the latter energy calculation can be conceptualised as a 3D convolution. Each convolution unit taps the spin shift register at appropriate points to feed a cubic kernel so that the values of nearest-neighbor spins arrive synchronously to be processed in parallel by the kernel as illustrated schematically in Fig. 1. The kernel unit handles the mathematical operations defined by the Hamiltonian by calculating the energy of a spin using the values

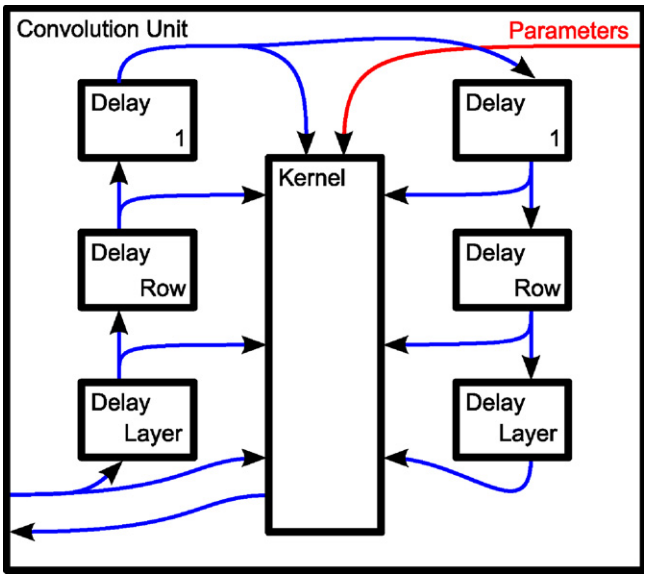

Fig. 1. The convolution unit employs delay lines to synchronously feed the data from the 6 nearest-neighbors to the kernel which processes the Hamiltonian energy calculation.

of cubic-nearest neighbor spins and the external field, as shown in Fig. 2. In this bitwise representation, the multiplications of neighboring spins is equivalent to the fundamental logic operation of eXclusive OR, of which the six resultants are then summed as either positive or negative one, and then summed with the bias field multiplied by the central cell. To avoid unnecessary calculation, we choose natural units with $J = -1$, and $b = \mu B_z$, so that the first term of the Hamiltonian gives an integer in the range $[-6, 6]$. Once the energy contribution of a spin in its current configuration (E) is known, a Heat-Bath algorithm is used to stochastically determine if the spin flips to the opposite state of energy (E'') using the probability:

$$P_{\text{flip}} = \frac{1}{1+\exp(E_\Delta/k_B T)} \tag{2}$$

where $E_\Delta = E'' - E$, i.e. the energy difference between the flipped (final) and unflipped (initial) calculated from Eq. (1). To reduce numerical calculations, we employ the standard result that:

$$E_\Delta = -2E$$

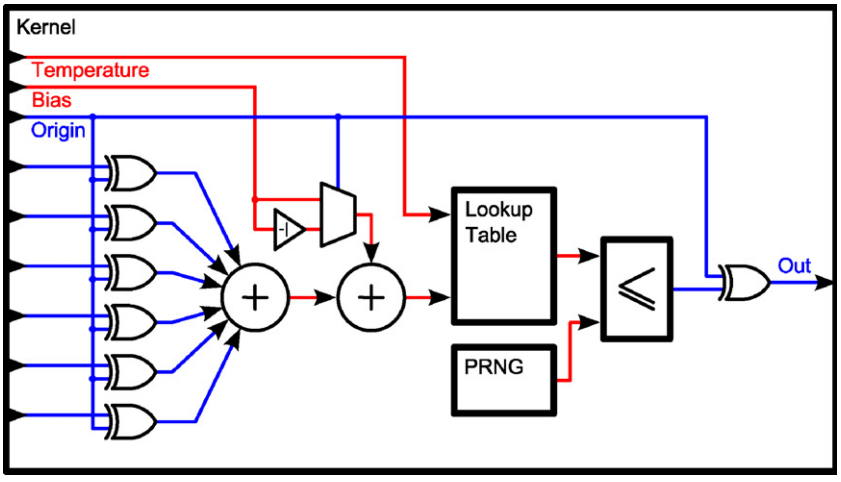

Fig. 2. The kernel performs the mathematical operations defined by the Hamiltonian by multiplaying the spin at the origin with the six nearest-neighbor values and adding the effect of an external magnetic field bias. The temperature parameter alters the probability for accepting a trial move which is read from a look-up table.

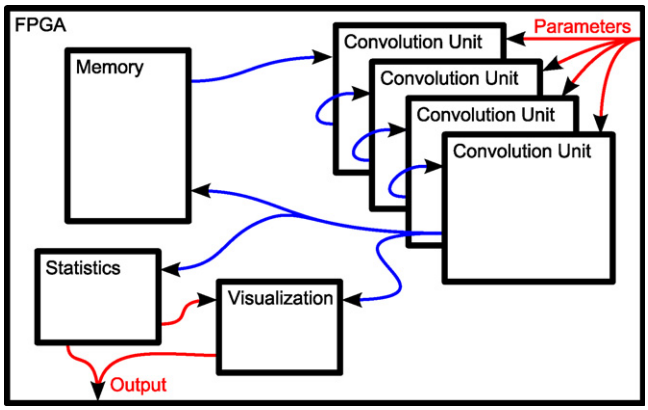

Fig. 3. The addition of convolution units allows for parallel calculation of spin updates at independent points in the array. Single bit wide paths are shown in blue and scalar paths are shown in red. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of the article.)

as $E' = -E$ which can be shown by letting $\sigma_i \rightarrow -\sigma_i$ in Eq. (1). This eliminates the need for 6 XOR gates and associated adders. Further optimisations are found in the probability function of the heat-bath algorithm. It comprises an exponential and division operation which would require significant resources to implement, but only depends on two variables: the energy shift (E_Δ), and the reduced temperature ($t = k_B T$) where k_B is the Boltzmann constant. These two variables are therefore absorbed into a look-up table resident in read only memory (ROM) which has pre-calculated values of the probability function for a discrete range of temperatures and fields. The spin-flip probability given by the ROM output is then compared against a Linear Feedback Shift Register which generates pseudo-random numbers in the range [0,1) with a period of 2^{32}. If the pseudo-random number is less or equal to the spin-flip probability, the flip is accepted, otherwise the spin remains in its initial state. Since, like the Metropolis algorithm, the heat-bath algorithm obeys the principal of detailed balance, the stochastic generation of new states eventually forces the ensemble to converge to the true statistical distribution despite the fact that the partition function is never explicitly calculated [12]. Defect cells representing dilute, non-magnetic impurities were included in our final program by duplicating the entire shift register loop with a second parallel set of shift registers storing the defect state (defect/not). The kernel required only minimal changes to incorporate this defect information, inhibiting the adder for each side of the cubic kernel that has a defect present on it. The completed kernel module is shown in Fig. 2. It is fed data from a string of shift registers as in the naive implementation of a convolution, with the output of the kernel being independent of the input allowing additional pipeline stages to be added if needed.

4. Parallelization

The notable thing about the FPGA implementation is the ease with which it can be parallelized simply by the addition of more convolution units, limited only by the available resources of the system. This can be viewed as a hardware analogue to a software approach that uses multi-threaded processes, where it is possible for each convolution unit/core to simultaneously perform calculations on independent sections of the spin lattice. This is schematically shown in Fig. 3. In this work we explored running between 1 and 24 convolution units simultaneously on our inexpensive system. The number of convolution units was limited by the spin lattice size: $32 \times 32 \times 32$ (24 cores); $64 \times 64 \times 64$ (8 cores); $128 \times 128 \times 128$ (2 cores); $512 \times 512 \times 128$ (1 core). By convention, a Monte Carlo step is defined as sufficient iterations of the main algorithm so that, on average, every spin in the array receives a single opportunity to alter its value. For this reason, the performance in Monte Carlo steps per second is inversely proportional to the volume (length) of the array. Larger lattices are therefore more costly. The embarrassingly parallel nature of the problem can be exploited to reduce the cost by chaining multiple convolution units, linearly increasing performance by the number of parallel processes. Each convolution unit operates at a constant rate of one spin update per clock cycle. Beyond our current implementation on a low cost development board with a single FPGA and external memory element, there are further opportunities for scaling the system both in array size and performance. The parallel implementation of the convolution units and their low bandwidth interconnect allows linear scaling of speed by adding additional devices/silicon. Our implementation operates with a 100 MHz clock while recently released devices surpass 400 MHz. Technical details regarding device resource usage are given in Fig. 9 in Appendix A.

4.1. Visualisation

Rather than having to share read access to the dataset in a single memory location, as in processor based implementations, both the visualisation and stastistics modules duplicate the memory write operations of the storage array to their own local memories. This allows them to process every spin of the lattice at each time step for implicit real-time presentation of the generated data without reducing the performance of the simulation, in contrast to other contemporary solutions [1]. To visualise the domains in the dilute antiferromagnet, it is insufficient to simply render the values of the spins at each site as this leads to a confusing image of many inter-penetrating checkerboards. Instead in this work we deploy a simple post-processing algorithm on the spin lattice image to distinguish between antiferromagnetic domains. The boundary between domains designates a point where two spins lie parallel to each other, separating larger clusters in which neighboring spins lie antiparallel to one aother. Past work has relied on unnecessarily complicated cluster methods to differentiate between antiferromagnetic domains in an Ising model [16]. In this work we used a simpler method to colour spins based on whether or not they obeyed the relation below:

$$\sigma_{x,y,z} = (-1)^{x+y+z}$$

where x, y and z are the Cartesian indexes of the spin in the cubic lattice. Green was used to colour spins which obeyed the above relation. Blue spins did not. Note, after this visualisation post-processing, solid blue (or green) regions actually contain Ising spins with both +1 and −1 values. Importantly, inside each of these regions antiferromagnetic ordering is unbroken so that perfect long-range ordering would correspond to a solid colour. Our method implicitly relies on the fact that only two types of domains (out of phase by 180°) are possible in a cubic Ising antiferromagnet. This fact can be easily checked conceptually by considering the translation properties of an infinite 2D or 3D checkerboard. Translation by a single grid point creates a distinct out-of-phase pattern, but translating a second grid point in any direction restores the initial phase. On an FPGA, this domain image processing was handled by the XOR of the spin with the least significant bits of its indices in the 3D array. Defect spins were rendered separately and represented as red pixels. Generally we examined only a 2D plane representing a single layer of the lattice.

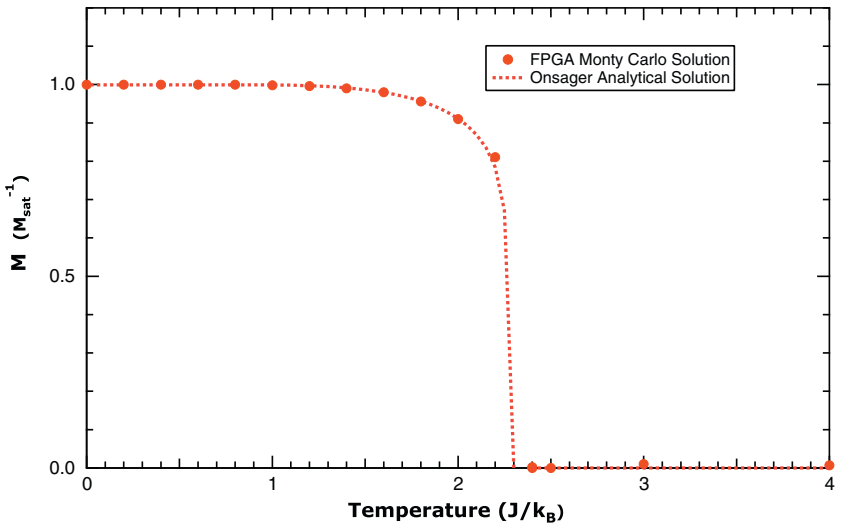

Fig. 4. The FGPA Monte Carlo numerical result for the average spin value (M) as a function of temperature in a 2D Ising ferromagnet agrees with the exact solution provided by Onsager. Computational simulation (circles) and analytical calculation (dotted line).

5. Results

5.1. Numerical accuracy: ordering temperature of undiluted 2D Ising model

A test of the numerical accuracy of the simulation was performed by comparing the solution against Onsager's analytical solution for a 2D Ising ferromagnet [17]. To facilitate direct comparison with the simpler case that can be exactly solved, a lattice size of $128 \times 128 \times 1$ with $J = 1$ and $b = 0$ and zero defects was simulated. Either fixed or periodic boundary conditions could be applied with no discernable effect on the interior ensemble averages. The system was placed in a fully saturated state at zero temperature, and then sequentially heated to higher temperatures allowing for 500 steps per temperature interval, and 5000 steps near the critical temperature to compensate for the slowing down of dynamics. After stability was reached at each temperature, the ensemble magnetisation was calculated from the average value of all N spins:

$$\langle M \rangle = \sum_{i}^{i \leq N} \frac{\sigma_i}{N} \tag{3}$$

Fig. 4 shows that the average $\langle M \rangle$ calculated from the FPGA simulation agrees well with the analytical solution giving a critical temperature $T_c \approx 2.4$. More precise analysis of numerical error in T_c requires the use of Binder accumulants [1] to avoid finite size effects, but this approach is of less interest here due to the obvious agreement for our larger simulation volumes. Smaller simulations, or simulations with insufficient MC steps, are well known to deviate from the exact result [18].

5.2. Meta-stable fractal domain state in bulk dilute antiferromagnet

The addition of random fields, due to disorder, in the Ising magnet results in the formation of frozen domains. This has been proven experimentally [5,9,13] and theoretically [3,15,16]. We used our simulation to investigate this feature. Following an argument similar to the original Imry-Ma argument [21], it can be predicted that domains should be sensitive to field cooling values but should be frozen at sufficiently low temperatures. We find that meta-stable domains are formed in the Ising simulation, as visualised in Fig. 5, which have a strong dependence on the degree of non-magnetic impurities and the magnitude of the cooling field. This behaviour is different from the larger domains that form in pure ferromagnet or antiferromagnet Ising magnets since these tend to be highly unstable and transient. Fig. 6 shows the surface layer of a 6 percent diluted DAFF field cooled from $t = 6.0$ to $t = 0.4$ for 500 MC's per 0.2 temperature interval in three different fields. Stronger external fields lead to the formation of a smaller and more chaotic domain structure [16]. Increasing dilution has a similar effect. The application of the zero or opposite field did not alter the domain distribution once the domains were formed, implying the structure is frozen at sufficiently low temperatures, and mirroring magnetic memory effects known to occur experimentally. Sufficiently high temperature removed the domain memory by forcing the ordered spins back into a random, paramagnetic

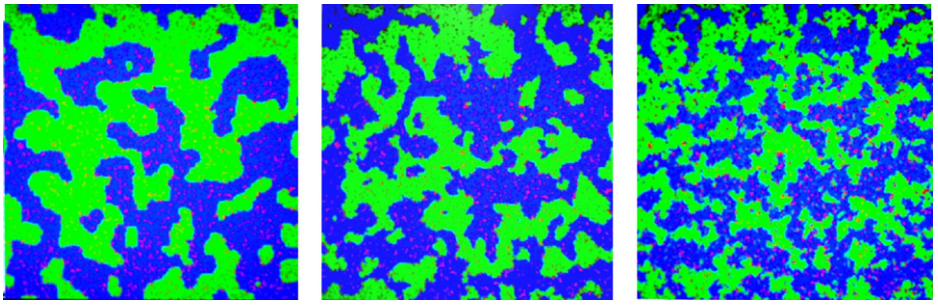

Fig. 5. A 2D plane extracted from the 3D lattice after post-processing shows sub-regions of antiferromagnetically ordered spins appear. The domain size depends on the magnitude of external field during cooling (left) $b = 0.5$ (center) $b = 2.0$ (right) $b = 4.0$.

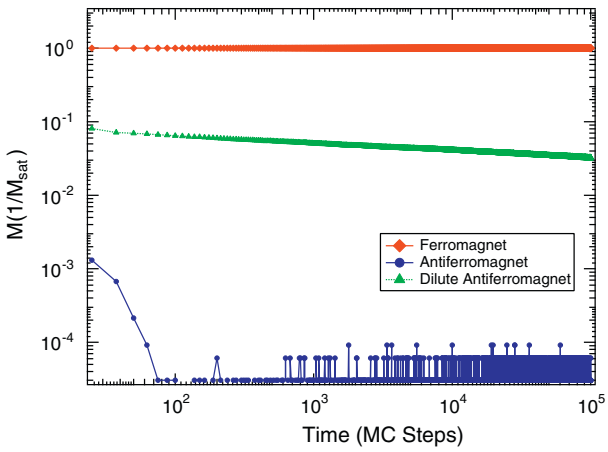

Fig. 6. After a large saturating field is suddenly turned off, the average value of spin behaves very differently in pure and dilute Ising magnets. For the ferromagnetic case, no relaxation occurs and a remnant magnetic state is preserved even in zero field which is the defining characteristic of a permanent magnet. For the antiferromagnet, the magnetisation quickly approaches zero as two antiparallel sub-lattices form providing perfect cancellation. In the dilute anti ferromagnet, such sub-lattice compensation is prevented by disorder so that a small parasitic magnetisation is preserved for long time periods and undergoes gradual non-exponential decay. Fits to the log–log relationship yield an exponent of $\alpha = 10.4$.

state. These features are in qualitative agreement with neutron diffraction experiments on $FeZn_xF_2 - x$ [5] which showed that antiferromagnet domain size was highly dependent on field cooling conditions and remained frozen below a critical temperature even when the field was removed. Based on these experimental and theoretical findings, an obvious point to make is that there must be a fundamental limit on the size of antiferromagnetic particle where a domain can form which is dependent on the field and dilution of the particle. This is discussed further in Section 5.4.

5.3. Altered relaxation Glauber dynamics in dilute antiferromagnet compared with ferromagnet

Although real spins have more complex dynamic behaviour than idealised Ising spins, Glauber pioneered a link between the Ising model and real dynamics showing that, in certain cases, the Ising model reveals insights into how systems approach equilibrium [8]. In this work, we use the FPGA implementation to show that the diluted antiferromagnet has highly distinct relaxation dynamics from the pure ferromagnetic or antiferromagnetic case. Each of the 3 different types of magnets was field cooled from $t = 6.0$ to $t = 0.4$, reducing temperature in 0.2 intervals while waiting 500 MC's per step, in an external field of $b = 2.0$. At this point the external field was switched off ($b = 0.0$) and data was collected regarding the system's reaction. Fig. 7 shows that the ferromagnet remained as a permanent magnet where local exchange forced spins to remain aligned. Meanwhile, spins in the antiferromagnet spontaneously and abruptly reordered into antiparallel alignments so that $\langle M \rangle$ was reduced rapidly to the level of thermal fluctuations in less than 100 MC steps. However, the diluted antiferromagnet showed a qualitatively different behaviour undergoing a rapid drop in magnetisation to about 10 percent polarisation followed by slow relaxation dynamics which follows a power-law shown by the straight line in the log–log plots. The fitted exponent is in good agreement with past theoretical findings [15]. This implies that the domains in the diluted antiferromagnet are capable of retaining some trapped magnetisation which is slowly destroyed by thermal activation leading to domain wall migration. This agrees qualitatively with the slow non-exponential decay law found experimentally [9,13].

5.4. Finite size effects: cross-over from nanoparticle to bulk behaviour

It has been noted by previous authors that domains do not form in small simulations [16], necessitating expensive large scale simulations, or unphysical field cooling values to create artificially small domains. This is noted even for simulations where periodic boundary conditions (mimicking an infinite lattice) are applied. Here we show that changing the boundary condition to a vacuum boundary (zero spin) accentuates this phenomenon so that domains are not formed for simulations of smaller than $64 \times 64 \times 64$ spins at a cooling field of $b = 0.5$ with 12.5 percent dilution. Although transient domains appear (as they do for the undiluted Ising model), these completely disappear between 100 and 5000 Monte Carlo steps resulting in single-domain states for smaller simulations in Fig. 7. Larger simulations begin to show the formation of domains for identical field and dilution values, and these domains persist well beyond 100,000 time steps with little change. The exact threshold point where a domain state can be frozen into the particle depends on three parameters: cooling field, dilution and size of simulation. We argue that the latter parameter is not merely a computational artifact but reflects a fundamental change in behaviour from nanocrystallites or nanoparticles to the case of bulk materials with vastly increased number of spins. If we assume a lattice constant of ≈0.5–1.0 nm which is typical for cubic Ising antiferromagnets (CoO or FeF_2) then, based on these preliminary results, we suggest that the critical grain size is between 30 and 60 nm for a multidomain state to form. More rigorous computational work based on more quantitative models would assist in determining the precise threshold points. However this qualitative feature is strongly suggested using the simple Ising model. Such a threshold would certainly explain puzzling discrepancies in various theories of magnetic exchange bias which disagree about whether domains play an important role [7,14].

6. Discussion and summary

We used a novel computational method to study the properties of the diluted Ising antiferromagnet. This class of material is fundamentally different from that of the pure ferromagnetic or antiferromagnetic version. The distinguishing features are the formation of a meta-stable domain structure and the possibility of a weakly decaying parasitic magnetisation in zero field. Notably domains only form for at least $64 \times 64 \times 64$ spin volumes (for $b = 0.5$

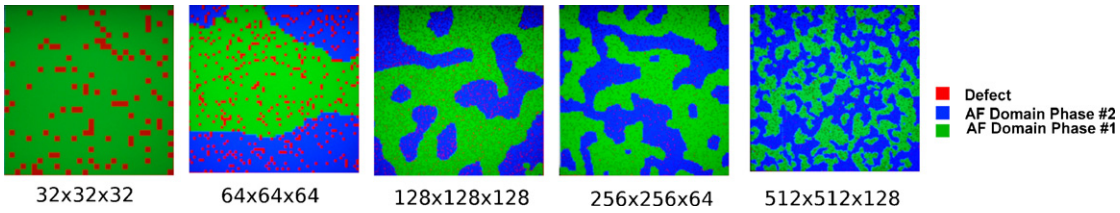

Fig. 7. Domain images for different lattices dimensions with 12.5 dilution and $b = 0.5$ after 1000 Monte Carlo steps, showing that multidomain states vanish for smaller simulation sizes. This suggests that domain behaviour is inhibited in antiferromagnetic nanoparticles or nanocrystalline solids beneath a critical size threshold.

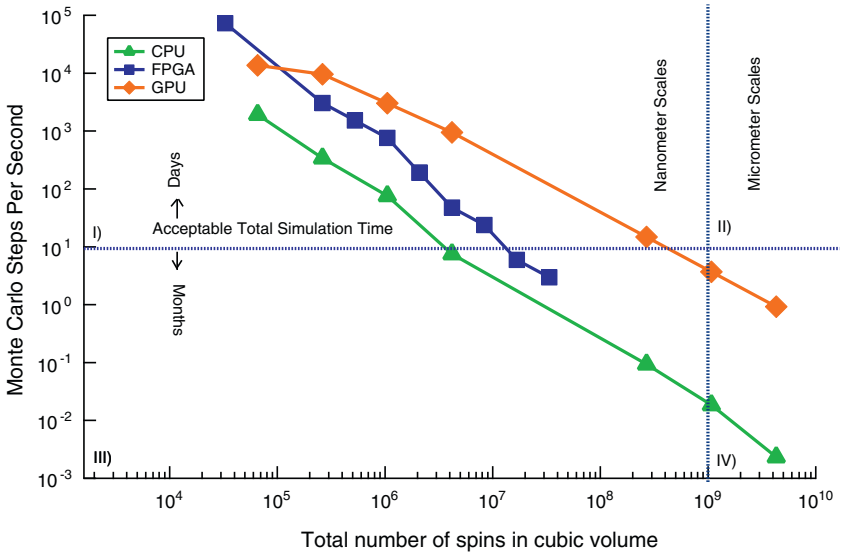

Fig. 8. Numerical benchmark of our single FPGA Ising model with contemporary single CPU and and single-die GPU. Data taken from Ref. [2]. FGPGA was Xilinx Spartan 3E 1200 with 100 MHz clock. CPU was a Intel Xeon X5560 at a clock rate of 2.80 GHz and 8192 kB cache. The 'multi-GPU' implementation was on a CUDA enabled NVIDIA Tesla C1060. All existing solutions fail to cross the threshold where a full micrometer simulation can be run in less than 7 days. To determine this threshold we assumed a complete simulation will typically need 100,000 Monte Carlo steps.

and low dilution of 12.5%) and so large scale atomistic simulations are essential to study these features. We tentatively suggest that the latter finite size effect is not merely a computational artifact but represents a fundamental change in behaviour from the nanoscale to the microscale. More rigorous theory and experiment should aim to certify this. Based on this preliminary result, if we estimate that the spins are 0.5–1 nm apart, the model predicts a critical length size where antiferromagnetic domains no longer form in the region around 30–60 nm. This threshold has technological consequences as this length scale is the typical grain size of nanocrystallites in polycrystalline thin films [7] suggesting that small grains may be incapable of forming multidomain states. The enhanced computational performance, and flexibility of the FPGA will enable future study of the effects of different geometric particle shapes, dilution and cooling field on this threshold point. Fig. 8 shows a performance comparison of our FPGA method with contemporary GPU and CPU implementations of the Ising model. Ostensibly the algorithm used on our single-chip FPGA is better for smaller systems ($\lesssim 64^3$) but does not scale as well as the GPU implementations. We remark that the exploratory work presented in this paper was performed on an older FPGA device (pre 2008), which, including the visual interface has less than 1Mb RAM. Extrapolating the performance figures to a target such as the SciEngines RIVYERA platform should place the developed algorithm to be competitive with the more expensive GPU systems even to much higher length scales. Fig. 8 also illustrates that all existing computational platform become unacceptably slow once micrometer scales are reached. Both a cluster of GPU's or a multi-chip FPGA supercomputer should soon be able to breach this limit, but the GPU version would consume orders of magnitude more power and space than an FPGA based system [11].

7. Conclusions

New computational techniques should aim to allow spin models to cross from the micrometer to the nanometer scales. The 2D and 3D Ising models were successfully implemented on a FPGA chip for the ferromagnetic, antiferromagnet and dilute antiferromagnetic cases. The computational performance assisted our study of distinct long time period behaviour of the three cases at large spatial scales approaching a micrometer. We showed that the FPGA approach was useful for running the large scale Monte Carlo simulations required to observe distinct features which do not appear for small simulations of a dilute antiferromagnet. FPGA implementations of the Ising model may have useful applications outside the field of magnetism. Further innovation in software, hardware and computational techniques is required to allow existing spin models to extend from the nanometer to micrometer range.

Acknowledgments

D.C. would like to express his sincere thanks for the financial support of the Australian Institute of Nuclear Science and Engineering(AINSE) and to Prof. R.L. Stamps and Dr. Klose for helpful discussions.

Appendix A.

$$H = -J\sum_{ij}\varepsilon_i\varepsilon_j\sigma_i\sigma_j - \sum_i \mu B_z \varepsilon_i \sigma_i$$

The first term describes Ising spins where σ_i is the value of a spin at a lattice point which can only take two values (+1 or −1).

```
Device utilization summary:
---------------------------
Selected Device : 3s1200efg320-4
Number of Slices: 4618 out of 8672 53%
Number of Slice Flip Flops: 640 out of 17344 3%
Number of 4 input LUTs: 9047 out of 17344 52%
Number used as logic: 836
Number used as Shift registers: 8211
Number of IOs: 85
Number of bonded IOBs: 83 out of 250 33%
Number of BRAMs: 5 out of 28 17%
Number of MULT18X18SIOs: 2 out of 28 7%
Number of GCLKs: 3 out of 24 12%
Number of DCMs: 1 out of 8 12%
Timing Summary:
---------------
Speed Grade: -4
Minimum period: 9.857ns (Maximum Frequency: 101.451MHz)
Minimum input arrival time before clock: 12.431ns
Maximum output required time after clock: 6.721ns
Maximum combinational path delay: 5.989ns
```

Fig. 9. FPGA hardware synthesis results for the $128 \times 128 \times 128$ spin lattice.

The binary nature of the Ising model represents a magnetic material with a large (infinite) crystalline anisotropy which forces spins to align along one Cartesian axis. This approximates highly uniaxial magnets such as CoO or FeF_2. Certain Ising spins are replaced by non-magnetic defects so that $\varepsilon_i = 0$ for a defect and $\varepsilon_i = 1$ for a magnetic spin. This means that Ising spins do not directly interact with defects, but the defects reduce the number of magnetic neighbors. The defect spins are randomly positioned and the total percentage of defect sites is termed the dilution ρ. The second term of the Hamiltonian describes the Zeeman coupling of individual spins with an externally applied magnetic field which tends to align them in the direction of the field. The magnetic moment per atom is described by μ. This term is responsible for paramagnetism above the critical temperature, and, in the case of the ferromagnetic Ising model, causes ferromagnetic domains to align with the external field below the critical temperature.

References

[1] K.A. Hawick, A. Leist, D.P. Playne, Interactive visualisation of spins and clusters in regular and small-world Ising models with CUD A on GPUs, J. Comput. Sci. 1 (2010) 33–40.
[2] B. Block, P. Virnau, T. Preis, Multi-GPU accelerated multi-spin Monte Carlo simulation of the 2D Ising model, Comput. Phys. Commun. 181 (2010) 1549–1556.
[3] D. Chowdhury, S. Kumar, Domain growth in the three dimensional dilute Ising model, J. Stat. Phys. 49 (1987) 855–858.
[4] J.H. Condon, A.T. Ogielski, Fast special purpose computer for Monte Carlo simulations in statistical physics, Rev. Sci. Instrum. 56 (1985) 1691–1696.
[5] A.R. King, D.P. Belanger, V. Jaccarino, Neutron scattering study of hysteresis near T_c in $d = 3$ random field Ising system, Solid State Commun. 54 (1985) 79–81.
[6] F. Belletti, et al., Simulating spin systems on IANUS, an FPGA-based computer, Comput. Phys. Commun. 178 (2008) 208–216.
[7] K. O' Grady, et al., A new paradigm for exchange bias in polycrystalline thin films, J. Magn. Magn. Mater. 322 (2010) 883–899.
[8] R.J. Glauber, Time-dependent statistics of the Ising model, J. Math. Phys. 4 (1963) 295–307.
[9] S.-J. Han, D.P. Belanger, Relaxation of the excess magnetization of random-field-induced metastable domains in FeZnF2, Phys. Rev. B 45 (1992) 972879–972881.
[10] N. Kazantseva, D. Hinzke, U. Nowak, R.W. Chantrell, U. Atxitia, O. Chubykalo-Fesenko, Towards multiscale modeling of magnetic materials: Simulations of FePt, Phys. Rev. B 77 (2008) 184428.
[11] S. Kestur, J.D. Davis, O. Williams, BLAS Comparison on FPGA, CPU and GPU 2010, IEEE Annu. Symp. VLSI 1 (2010) 883–899.
[12] D. Loison, C.L. Qin, K.D. Schotte, X.F. Jin, Canonical local algorithms for spin systems: heat bath and Hasting's methods, Eur. Phys. J. B 41 (2004) 395–405.
[13] R. Bruinsma, J. Hammann, M. Lederman, J.V. Selinger, R. Orbach, Low-temperature dynamics of a diluted Ising antiferromagnet, Phys. Rev. Lett. 68 (1992) 2086–2089.
[14] P. Miltenyi, M. Gierlings, J. Keller, B. Beschoten, G. Guntherodt, Diluted antiferromagnets in exchange bias: proof of the domain state model, Phys. Rev. Lett. 84 (2000) 0031–9007.
[15] U. Nowak, K.D. Usadel, Nonexponential relaxation of diluted antiferromagnets, Phys. Rev. B 43 (1991) 851–853.
[16] U. Nowak, K.D. Usadel, Structure of domains in random Ising magnets, Phys. Rev. B 46 (1992) 8329–8335.
[17] L. Onsager, Crystal Statistics: a two-dimensional model with an order-disorder transition, Phys. Rev. 65 (1944) 117–149.
[18] T. Preis, P. Virnau, W. Paul, J.J. Schneider, GPU accelerated Monte Carlo simulation of the 2D and 3D Ising model, Comput. Phys. Commun. 228 (2009) 44684477.
[19] R.L. Stamps, Dynamic magnetic properties of films, multilayers and patterned elements, Adv. Funct. Mater. 20 (2010) 2380–2394.
[20] T. Guneysu, C. Paar, J. Pelzl, G. Pfeiffer, M. Schimmler, C. Schleif-fer, Parallel computing with low-cost FPGA's: a framework for COPA-COBANA, Adv. Parallel Comput. 15 (2008) 741–748.
[21] U. Nowak, K.D. Usadel, Domain state model for EB. I. theory, Phys. Rev. B. 66 (2002) 014430.

David Cortie is a PhD student with the Australian Nuclear Science and Technology Organization and the University of Wollongong, studying the physics of future spintronic materials. His research interests include computational physics, frustrated interactions in condensed matter and nanomagnetic systems studied with polarised neutron scattering.

An early career professional engineer working across diverse areas in multidisciplinary teams, **John Pillans** brings a broad background of experience to his role as systems engineer. Having taught digital design at the university level, he brings a contemporary knowledge to FPGA projects, and is an evangelist for novel applications of the technology. He is currently in a research role at a prominent multinational technology and will be innovating in less visible environments.