Design Tradeoffs for Improved Performance in MQCA-based Systems

Michael Niemier, Aaron Dingler, and X. Sharon Hu
Department of Computer Science and Engineering, University of Notre Dame
Notre Dame, IN 46556, USA
Email: {mniemier,adingler,shu}@nd.edu

Abstract—**This paper looks at circuits and systems made from nanoscale magnets. We leverage physical level simulation to show how to extend initial experimental work – where wires and gates with nanometer feature sizes have been demonstrated at room temperature – to show how one might make design tradeoffs to improve both the performance and reliability of more complex systems.**

I. INTRODUCTION

At present, numerous research efforts are looking for a new logic device either to replace or augment CMOS technology in order to continue the performance scaling trends that we have seen for the last 40 years. This paper looks at how a magnetic implementation of the Quantum-dot Cellular Automata (QCA) device architecture might help in this regard. We will first discuss techniques that we have developed to analyze nanomagnet-based QCA (MQCA) circuits. Then, we will show how such analysis can be used to determine whether or not systems of MQCA devices might offer performance wins over the projected state of the art in CMOS for more systems-level tasks.

QCA accomplishes logical operations and moves data via nearest-neighbor interactions rather than with electric current flow. Initial experiments considered a *metal-dot* implementation of a QCA device [1]. A theoretical analysis of a *molecular* implementation claims that QCA-based circuits could be clocked at an extremely high frequency, lead to circuits with densities that are 3-4 orders of magnitude beyond what end-of-the-curve CMOS can provide, and should dissipate very little power [2]. Experiments indicate that a molecule could switch between configurations that could represent binary states [3]. *Semiconductor-based* QCA devices have also been realized [4] and could potentially play a role in quantum computers. This work focuses on *magnetic* QCA [5] – what we believe to be the most promising and practical implementation of this device architecture for at least the near- to mid-term[1].

II. BACKGROUND

A. QCA Circuit Constructs

The initial description of a QCA device called for encoding binary numbers into cells that have a bi-stable charge configuration. QCA cells interact because the charge configuration of one cell alters the charge configuration of the next cell. In a magnetic implementation of QCA, charge configurations are replaced with magnetic polarizations of single domain magnets. Fig. 1 illustrates two important building blocks that would be used to construct QCA circuits. A wire (Fig. 1a) is just a line of magnets that are antiferromagnetically coupled with each other. The basic logic gate in MQCA is based on the majority voting function – where the output is the logical value associated with the majority of the 3 input cells. By setting one

---

[1]Note that MQCA is used in the realm of classical computing and should not be mistaken for a quantum computing architecture.
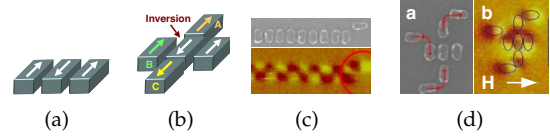

Fig. 1. Cartoon representations of (a) a wire segment and (b) a majority gate. Wire segments have been experimentally demonstrated (c) as have majority gates (d). Crossovers function correctly in simulation.
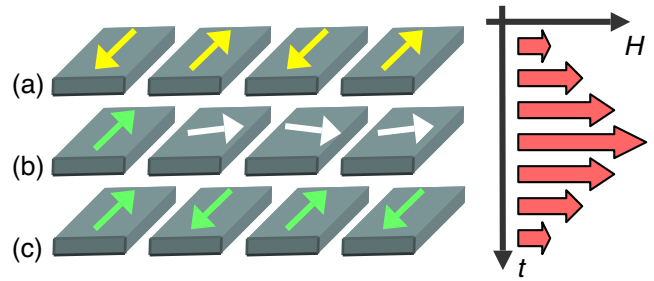

Fig. 2. Operating scheme of a wire: (a) initial configuration, (b) high-field ("null") state, (c) after the application of the input, and the final ordered state.

input of a majority gate to a logic '0' or '1', the gate will execute an AND or OR function respectively. Note that in MQCA, the gate performs an *inverting* majority voting function (Fig. 1b).

B. Required Clock Functionality

The structures illustrated in Fig. 1c,d were tested with a clock that took the form of a periodically oscillating external magnetic field that drove a system to an initial state, and then controlled the relaxation of the said system to a ground state. Conceptually, a magnetic field is applied along the hard (or shorter) axes of a group of nanomagnets. For example, Fig. 2a, illustrates a line (or "wire") of nanomagnets that has relaxed to a logically correct, antiferromagnetically coupled ground state. In Fig. 2b, the external field turns the magnetic moments of all magnets horizontally into a neutral logic state against the preferred magnetic anisotropy (i.e. along the hard axes of the magnets). This is an unstable state of the system, and when the field is removed, the nanomagnets relax into a new antiferromagnetically ordered ground state in accordance with the new input (Fig. 2c.) If the first dot of the chain is influenced by an input device during relaxation, then its induced switching sets the state of the whole chain due to the bipolar coupling. [6] explored the use of copper wires wrapped with ferromagnetic material on the sides and bottom to provide on-chip, local control of MQCA-based circuits. [6] notes that field strengths of approximately 4,000 A/m to approximately 40,000 A/m could be produced in a 2 μm wide, 4 μm long and 200 nm thick wire yoked in ferromagnetic material with current densities of 10^6 A/cm^2 and 10^7 A/cm^2 respectively – with stronger fields obviously requiring more energy.

C. Scalability and Potential Performance Wins of MQCA

When considering MQCA as a candidate technology to extend the performance scaling trends associated with CMOS, MQCA has many desirable features. For MQCA, wires, gates, and inverters have all been experimentally demonstrated and operate at **room temperature** (see Fig. 1c,d [5]). MQCA also has the potential to offer **true low power operation**. [7] estimates that if 10^{10} magnets switch 10^8 times/second, they would only dissipate about 0.1 W of power. Moreover, if the drive circuitry is considered, [6] predicts that an MQCA-based circuit could provide performance wins over state-of-the-art, low power CMOS when considering the important energy-delay product (EDP) benchmark.

MQCA-based circuits are **scalable**. Nanomagnets with feature sizes on the order of 20 nm should still be above the superparamagnetic limit and are also non-volatile (although shape is also a factor). Additionally, [8] notes that as device sizes decrease below the superparamagnetic limit, binary state should be stable for 1 ms – enough time for the nanomagnets to be used to process information. Scaling can also decrease switching times [8].

Fabricating MQCA-based circuits and systems should be fairly straightforward. Nanomagnets can be made using conventional lithography or by leveraging electron beam lithography and lift-off to form specific patterns of magnetic material. One could also leverage imprint lithography where imprint is used to make molds, and then the molds can be used to make the nanomagnet shapes. Moreover, at least a subset of these methods is compatible with and can take advantage of advances in CMOS fabrication techniques – suggesting the possibility of heterogeneous technology integration. By leveraging fabrication processes that are well understood, manufacturing variations and their associated **defects should be more readily controlled**. We note that with MQCA-based circuits, there could be ways to eliminate faulty behavior even post-fabrication. Finally, with MQCA, **designing the other structures required to build more complex circuits and systems (wire crossings, fanout, etc.) appears to be quite possible to do.**

III. DESIGNING AND ANALYZING CIRCUITS

In this section, we introduce a methodology to design and analyze more complex circuit structures. More specifically, we consider how both external structures (the clock) and physical structures (the magnets) can affect the removal of the binary state associated with a previous computation. This is essential as it allows the nanomagnets that make up MQCA circuit elements to be reevaluated with new inputs. Understanding the interactions between these two design elements – the external magnetic clocking field and the nanomagnets themselves – is especially important. It will allow us to design circuit elements that not only perform tasks of interest (e.g. fanout, wire crossings, etc.), but also said structures should perform better (e.g. a wire that might be nulled with a clock that requires less energy).

We will explain how interactions between (a) an external stimulus such as the clock and (b) the nanomagnets can affect both performance and functionality by discussing four micromagnetic simulations of supermalloy nanomagnets performed with the OOMMF simulation suite [9] developed by NIST. As in [10], we assumed a saturation magnetization of 8.0×10^5 A/m, an exchange stiffness constant of 1.05×10^{-11} J/m, an anisotropy constant K_1 of 3 J/m^3, and the default damping constant of 0.5. Each simulation stage (when the magnitude or direction of the applied field changes) was considered complete when the maximum $|d\mathbf{m}/dt|$ dropped below a preset number
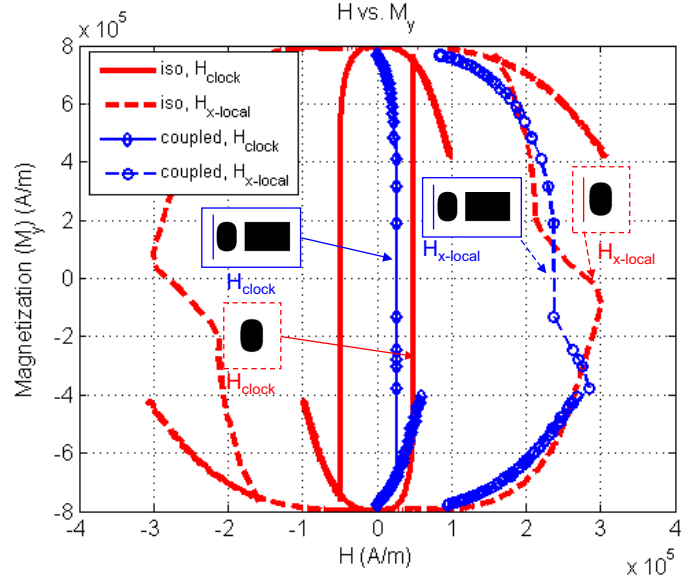

Fig. 3. The results of four different simulations are shown. The two simulations represented by solid line curves illustrate the external field required to null a 60x90x30 nm magnet both in isolation and coupled with a block of magnetic material (magnetized in the direction of the applied field). The magnet coupled with magnetic material requires a lower external field to switch because of the contribution of the nulling field from the surrounding magnetic material. The dashed line curves illustrate the field seen near the surface of the magnet. These local fields are similar when the 60x90x30 nm magnet is nulled (i.e. $M_y = 0$).

of degrees per nanosecond. We note that OOMMF is widely used and there is excellent correlation between simulation and experimental results (see [11], [12] for examples). All results are summarized in Fig. 3.

We begin by discussing a simulation where we consider a 60x90x30 nm magnet initially pointed down. We then applied an external magnetic field orthogonally to the magnet (i.e. along its hard axis) as well as a 15,000 A/m biasing field along its easy axis. The biasing field is necessary as, after we remove the y-component of the magnetization with the external nulling field, we want to "tip" it to the opposite polarization. Without a biasing field, the magnet would randomly tip up or down.

In more detail, the curves shown in Fig. 3 illustrate how the y-component of magnetization changes as a function of field strength. Consider the simulation "iso, H_{clock}" in Fig. 3. This 60x90x30 nm magnet in isolation is initially pointing down and the magnitude of the y-component is strongly negative. As the applied external field (H_{clock}) is increased, the y-component of magnetization is removed from the nanomagnet – eventually reaching 0 when the external field reaches a value of approximately 0.5×10^5 A/m. In the context of our discussion, the nanomagnet is now *nulled* and is tipped up by the biasing field. Note that if H_{clock} continues to increase, it will work to keep the magnet nulled. This explains why after the negative component of M_y is completely removed, that a strong, positive M_y component is not immediately seen (see the "tails" on the bottom-left and top-right of the "iso, H_{clock}" curve). Thus, the field can be too strong. When the nulling field is eventually removed, the y-component of magnetization is strongly positive. After the magnet has changed polarizations, we again apply an external field along the magnet's hard axis and a biasing field of the same magnitude but in the opposite direction to return the magnet to its original state.

(We repeated these simulations with an H_{bias} of different

magnitudes and noted a relationship between the field required to null the magnet and the field that biases the magnet into its new state. A stronger biasing field generally meant that the magnitude of H_{clock} required to move M_y to 0 would decrease. Note that in these simulations, the bias is externally generated, but in an MQCA-based circuit, it would come from the nanomagnets that actually comprise the circuit – which should ensure correct antiferromagnetic coupling as seen in Fig. 1(d).)

By determining where $M_y = 0$, one can determine the magnitude of the external field required to null a given magnet. Again, for these three simulations, the external field required was about 0.5×10^5 A/m. This is actually quite high and, if truly necessary, would have a significant impact on the energy required to perform a computation with nanomagnets. Leveraging work in [6], we determined that a current density of more than 10^7 A/cm^2 would be required to generate this clocking field. This will not only dramatically increase the energy of the MQCA system, but heating and electromigration effects would also be a concern.

However, it is important to realize that magnetic material also contributes to the *local* nulling field that a magnet sees [6]. To illustrate, we again consider a 60x90x30 nm magnet, but this time add a block of magnetic material that is polarized in the direction of the external clock ("coupled, H_{clock}"). Using the same 15,000 A/m bias, the 60x90x30 nm nanomagnet in this system was nulled when H_{clock} reached 0.25×10^5 A/m – about 2x lower than the field required to null the magnet in isolation. We also show the relationship between the x-component of the field seen at the *surface* of the nanomagnet and its actual magnetization. This value ($H_{x-local}$) is more insightful as it also incorporates the contribution of the magnetic material itself. One way to interpret this data is that in order to remove a magnet's remanent magnetization and reuse it to evaluate new inputs, it needs to "see" this value of $H_{x-local}$. Depending on the size and shape of the surrounding magnetic material, a given value of $H_{x-local}$ may be generated with a higher or lower *external* field.

This is best illustrated by considering the first simulation discussed in this section – the 60x90x30 nm magnet that is nulled in isolation (see "iso, $H_{x-local}$"). Note that the magnitude of $H_{x-local}$ in this system is nearly identical to that of the system with the block when $M_y = 0$. However, the system with the block requires a lower external clocking field in the presence of the same bias.

From a circuit and systems standpoint, one can infer that the system with the block requires less energy from the clock. Thus, when an external clock is applied to null the magnets, flux leaving the magnetic material that comprises various circuit structures can actually assist in the nulling process. We can leverage this idea not only to design gates and wires that operate at lower energy (i.e. by using magnetic material to increase the local bias that a magnet sees), but also to design new, fault-resistant circuit structures.

IV. DESIGN TRADEOFFS

We can leverage the work discussed in Sec. III to explain and consider performance tradeoffs. We offer three examples.

A. Magnitude of Applied Clock Field

As we have seen, the strength of the external magnetic field is important in facilitating the removal of remanent magnetization as it allows magnets to be reused to evaluate new logical inputs. We now leverage the above design techniques (a) to consider the behavior of an MQCA circuit element – a short wire segment, (b) to show why faulty behavior might occur,

and (c) to discuss briefly systems-level design tradeoffs that one might leverage to eliminate faulty behavior.

More specifically, we consider logical correctness in a wire as a function of field strength. We first discuss the results of two OOMMF simulations of a single wire that consisted of five, 60x90x30 nm magnets, spaced 15 nm apart, that were driven by an external field with a 2 ns period[2]. Thus, the field was applied to this system for 1 ns, turned off, and the magnets were allowed to relax to a new ground state for 1 ns. The last magnet in the wire was coupled to a block of magnetic material as illustrated in Fig. 3. To initialize the simulations, the wire is forced into an antiferromagnetically aligned state and the first magnet's polarization is manually flipped so it is ferromagnetically aligned with the second magnet. This simulates a new input to the wire. The external field is then applied and the magnets are expected to realign according to the new input (see Fig. 2).

The dependence on the magnitude of the external field is clear: when a 3,183 A/m (4 mT) field is applied, the new input does not propagate down the wire correctly. However, when a 6,366 A/m (8 mT) field is applied, the magnets correctly align (see Fig. 4 insets). We can use the observations discussed in Sec. III to explain. In the 3,183 A/m simulation, the fourth cell in the wire remains polarized along its hard axis, which incorrectly influences the fifth cell to trend up. Since the input is down, the fifth cell should also be down at the end of the simulation. When the field is raised to 6,366 A/m the five cells are all antiferromagnetically coupled in the order specified by the input. Why the higher field magnitude causes the input to propagate correctly is made clear by examining the local field between the third and fourth cells. Fig. 4 shows how this field changes over the course of the simulation and how it affects the magnetization of the fourth MQCA cell. The magnitude of the local field for the 3,183 A/m simulation does not reach the same peak as in the 6,366 A/m simulation, and the magnet does not switch. The $H_{x-local}$ associated with the 6,366 A/m simulation is required for the fourth magnet to polarize correctly. A 3,183 A/m applied field is simply not enough to generate this peak magnitude between the third and fourth magnets and causes a stuck-at fault.

Here, faulty behavior is eliminated by increasing the energy supplied to the system – i.e. by raising the external field generated by a clock. That said, other tradeoffs and design decisions might be made to spare this increase in energy.

B. Spacing Between Magnets

Spacing between adjacent magnets also affects the overall functionality of an MQCA circuit. Intuitively this makes sense: cells that are closer together will be more tightly coupled. As the magnets are moved farther apart, they have less influence on each other during the nulling process. An experiment similar to that discussed in Sec. IV-A demonstrates the impact of cell spacing on the functionality of MQCA wires. We again simulate a 5 cell wire – with H_{clock} initially equal to 3,183 A/m. When the magnets are separated by 12 nm, all magnets in the line switch correctly in accordance with the new input. However, if the separation is increased to 15 nm, the fourth magnet in the line fails to switch correctly given this value of H_{clock}. Fig. 5a plots M_y of the fourth magnet versus $H_{x-local}$ between the third and fourth magnets for each simulation. For each trace, $H_{x-local}$ is sampled 6 nm away from the fourth magnet. This graph emphasizes that, for a given size and shape of magnetic material, we need to generate a local magnetic field of a given magnitude to facilitate switching. For the successful

---

[2]A block of magnetic material like that shown in Fig. 3 was used to manage the boundary condition of the last magnet.
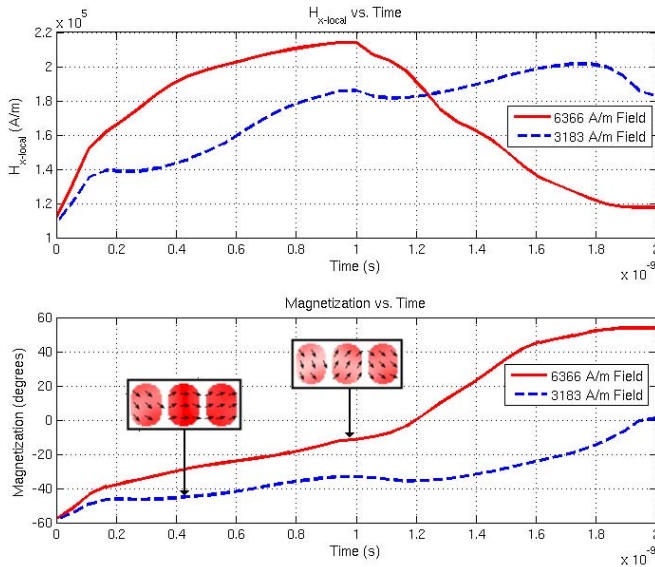
Fig. 4. $H_{x-local}$ and M_y vs. time for two different values of H_{clock}. In the simulation with faulty behavior, the magnitude of H_{clock} is not sufficient to generate the required local magnetic field and its magnetization does not flip.
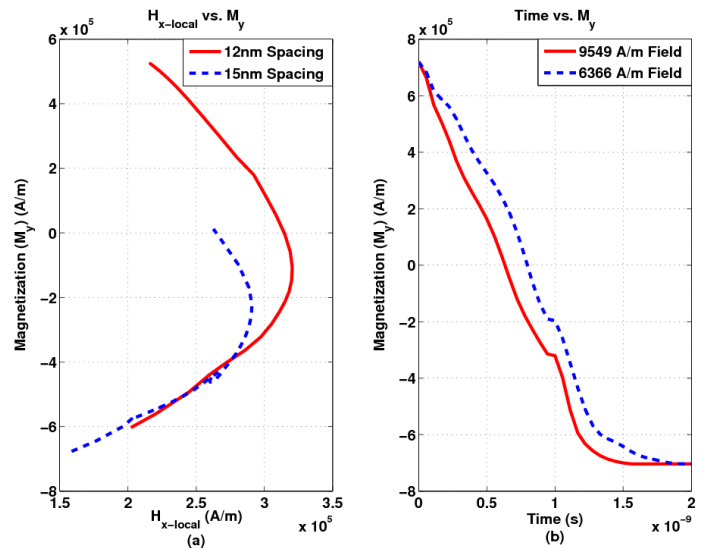

Fig. 5. (a) As the spacing between magnets in a line increases, faulty behavior can ensue as a given magnet is not nulled. For the faulty case, the magnitude of $H_{x-local}$ is too low. (b) When the magnitude of H_{clock} increases, a magnet can be nulled and biased toward the correct logical state more quickly.

simulation (12 nm spacing), the y-component of magnetization for the fourth magnet transitions from strongly negative to strongly positive. $M_y = 0$ when $H_{x-local}$ is approximately 3.2×10^5 A/m. For the flustered system, $H_{x-local}$ never reaches the same magnitude and a negative to positive transition for M_y is not seen. From a systems-level perspective, this result suggests that if H_{clock} is fixed (perhaps because of an overall energy budget) greater manufacturing precision can help to ensure logical correctness. We can leverage similar simulations to help identify required manufacturing tolerances if a target energy budget is known.

C. Field Strength and Time

If energy is less of a concern, and the magnitude of H_{clock} can be increased, switching time can *decrease*. For a line of 60x90x30 nm magnets, we see from Fig. 5b that, with a stronger clock, a given magnet can be nulled and correctly biased more quickly. This is important as it can in turn bias the next magnet more quickly, decreasing total time for a line to switch completely.

V. CONCLUSIONS

In conclusion, we have illustrated advantageous features of a magnetic implementation of QCA, presented a methodology for quantifying the functionality of MQCA devices, and leveraged that methodology to show how one might make performance tradeoffs – with an end goal of showing performance improvements over the projected state of the art for CMOS. Application spaces should be abundant as the devices should be low power and non-volatile, and any application that has these performance requirements might benefit. Patterned thin-film nanomagnets are also similar in nature and compatible with the processing and fabrication of MRAM devices – which among other things suggests another way to implement a processing-in-memory architecture. We should also be able to capitalize on advances made in magnetic data technology to address input and output in MQCA systems. Future work will leverage the above methodologies to design magnets to perform other circuit-level tasks to facilitate such applications.

REFERENCES

[1] I. Amlani, A. Orlov, G. Toth, G. Bernstein, C. Lent, and G. Snider, "Digital logic gate using quantum-dot cellular automata," *Science*, vol. 284, p. 1999, 289-291.
[2] J. Timler and C. Lent, "Power gain and dissipation in quantum-dot cellular automata," *J. of App. Phys.*, vol. 91, pp. 823–831, 2002.
[3] H. Qi, S. Sharma, Z. Li, G. Snider, A. Orlov, C. Lent, and T. Fehlner, "Molecular quantum cellular automata cells. electric field driven switching of a silicon surface bound array of vertically oriented two-dot molecular quantum cellular automata," *J. Am. Chem. Soc.*, vol. 125, pp. 15 250–15 259, 2003.
[4] L. Hollenberg, A. Greentree, A. Fowler, and C. Wellard, "Two-dimensional architectures for donor-based quantum computing," *Physical Review B*, vol. 74, pp. –, 2006.
[5] A. Imre, G. Csaba, L. Ji, A. Orlov, G. Bernstein, and W. Porod, "Majority logic gate for Magnetic Quantum-Dot Cellular Automata," *Science*, vol. 311, No. 5758, pp. 205–208, January 13, 2006.
[6] M. Niemier, M. Alam, X. Hu, G. Bernstein, W. Porod, M. Putney, and J. DeAngelis, "Clocking structures and power analysis for nanomagnet-based logic devices," *Proc. of Int. Symp. on Low Power Elec. and Des. (ISLPED)*, pp. 26–31, 2007.
[7] G. Csaba, P. Lugli, A. Csurgay, and W. Porod, "Simulation of power gain and dissipation in field-coupled nanomagnet," *J. of Comp. Electronics*, vol. 4(1/2), pp. 105–110, 2005.
[8] X. Wu, C. Liu, L.Li, P. Jones, R. Chantrell, and D. Weller, "Non-magnetic shell in surfactant-coated FePt nanoparticles," *J. Appl. Phys.*, vol. 95, pp. 6810–6812, 2004.
[9] M. Donahue and D. Porter, "OOMMF User's Guide, Version 1.0, Interagency Report NISTIR 6367," http://math.nist.gov/oommf.
[10] N. Dao, S. Whittenburg, and R. Cowburn, "Micromagnetics simulation of deep-submicron supermalloy disks," *J.of Appl. Phys.*, vol. 90(10), pp. 5235–7, 2001.
[11] L. Verma and V. Ng, "Magnetic domain patterns in a zigzag nanowire," *J. of Magnetism and Magnetic Materials*, vol. 313(2), pp. 317–321, 2007.
[12] S. McVitie, G. White, J.Scott, P. Warin, and J. Chapman, "Quantitative imaging of magnetic domain walls in thin ?lms using lorentz and magnetic force microscopies," *J. of Appl. Phys.*, vol. 90(10), pp. 5220–7, 2001.