Apprenticeship at the IBM Almaden Research Center

Jan, 03-Dec, 03

Ling Chen

From Jan 6, 2004 to Dec 19, 2004, as part of the requirement of our Ph.D. program, I worked as an apprentice under manager Barbara Jones at the IBM Almaden Research center.

About IBM

International Business Machines (IBM) is the world's top provider of computer hardware. Among the leaders in almost every market in which it competes, the company makes desktop and notebook PCs, mainframes and servers, storage systems, and peripherals, among its thousands of products. The company's service arm is the largest in the world. IBM is also one of the largest providers of both software and hardware.

Though best known for its hardware, IBM is shifting its focus primarily to the IT services business. IBM acquired PricewaterhouseCoopers' consulting and IT services unit, PWC Consulting, for an estimated $3.5 billion in cash and stock. This transaction serves the dual purpose of augmenting IBM's standard array of outsourcing, maintenance, and integration services, while moving the company into high-end management consulting[1]. The acquisition also helps IBM push its On Demand strategy, which basically means companies can buy computing needs from IBM like buying electricity from the power company.

About IBM Almaden Research Center

IBM’s Almaden Research Center (ARC) is one of eight laboratories worldwide. Located in the Santa Clara County Park southeast of San Jose, this laboratory occupies 690 acres in the hills above Silicon Valley.

Because of the huge loss of its hard disk department, IBM sold 70% of the hard disk

department to Hitachi. Now a whole wing of the Almaden Research center belongs to Hitachi. And many former IBM employees now work for Hitachi.

Almaden's approximately 400 permanent staff are doing research in two primary areas: advancing fundamental scientific understanding and working with IBM product developers to use this knowledge to create and improve IBM products.

Within IBM, Almaden is the heart of database research (Almaden researchers invented the relational database in the 1970s). It also has strong research in physical and materials science and storage systems. The Giant Magnito Resistance technology was developed in ARC. My work at ARC was mainly about understanding the physical phenomena that may advance the Magnetic Random Access Memory technology.

Non-Volatile Magnetic Random Access Memory

A dynamic random access memory (DRAM) retains its memory even after removing power from the device. This non-volatile memory has important applications in many areas. Such a device could have important commercial applications if the non-volatility were achieved without compromising other properties of the memory, such as density, read and write speed, and lifetime. IBM has recently begun a project to study the feasibility of a DRAM memory using memory cells based on magnetic tunnel junctions.

Magnetic random access memory (MRAM) (figure 1) stores data in the spin of electrons inside tiny magnetic sandwiches[2]. The sandwiches don't lose their magnetism—or memory—even if the power goes out. What's more, MRAM has the potential to store as much data as DRAM, write it faster and access it almost instantly—all while consuming less energy. A magnetic random-access memory cell is made of a thin insulating layer sandwiched between two magnets. Depending on the way their electrons are spinning, the magnets can have fields that point in the same direction or in opposite directions. These two states correspond to the ones and zeroes of digital memory. Each cell sits at the intersection of perpendicular electrodes that run above and below it. To write a bit,

electric current must pass through two intersecting electrodes. At the cross point, the currents induce a magnetic field sufficient to alter the spins of the electrons and therefore the cell's magnetic orientation. To read the data, a lower-voltage current travels along the bottom electrode, through the specific memory cell, and out along the top electrode. If the two layers are magnetized in the same direction, their electrical resistance is low, indicating a one; if they are magnetized in opposite ways, the resistance is high, indicating a zero. An electrical sensor at the end of the top electrode reads the resistance and determines the binary state of the cell that was read. Motorola, IBM and Eden Prairie are currently developing this technology.

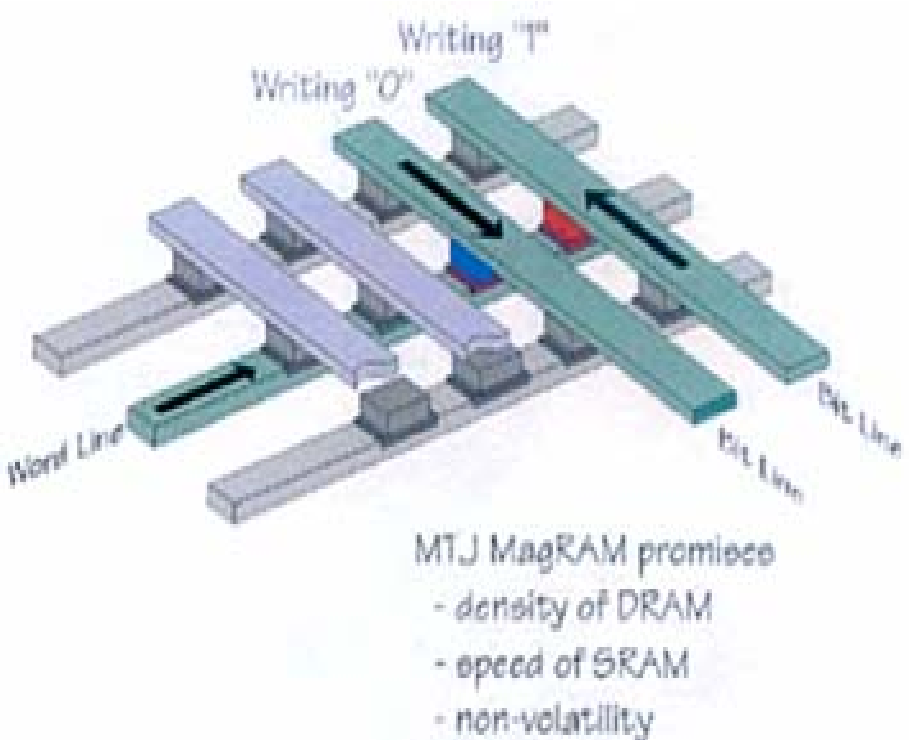

Figure 1. Magnetic Random Access Memory (MRAM)

Comparison with three major categories of memories available in the market[3]

1. Dynamic Access Memory

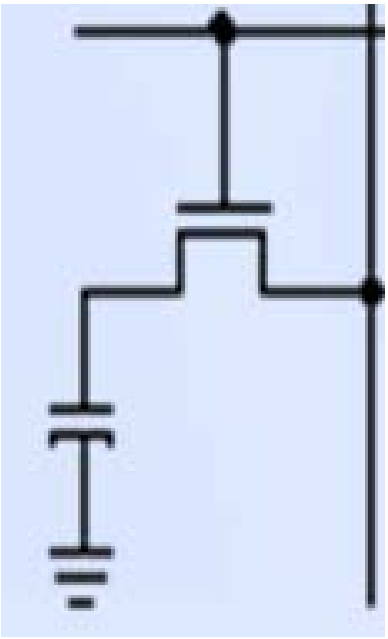

Figure 2. A single DRAM cell consists of a capacitor and a transistor.

Dynamic RAM is the most common type of memory in use today. Inside a dynamic RAM chip, each memory cell holds one bit of information and is made up of two parts: a transistor and a capacitor (figure 2). These are, of course, extremely small transistors and capacitors so that millions of them can fit on a single memory chip. The capacitor holds the bit of information -- a 0 or a 1. The transistor acts as a switch that lets the control circuitry on the memory chip read the capacitor or change its state.

A capacitor is like a small bucket that is able to store electrons. To store a 1 in the memory cell, the bucket is filled with electrons. To store a 0, it is emptied. The problem with the capacitor's bucket is that it has a leak. In a matter of a few milliseconds a full bucket becomes empty. Therefore, for dynamic memory to work, either the CPU or the memory controller has to come along and recharge all of the capacitors holding a 1 before they discharge. To do this, the memory controller reads the memory and then writes it right back. This refresh operation happens automatically thousands of times per second.

This refresh operation is where dynamic RAM gets its name. Dynamic RAM has to be dynamically refreshed all of the time or it forgets what it is holding. The downside of all of this refreshing is that it takes time and slows down the memory.

2. Static Memory

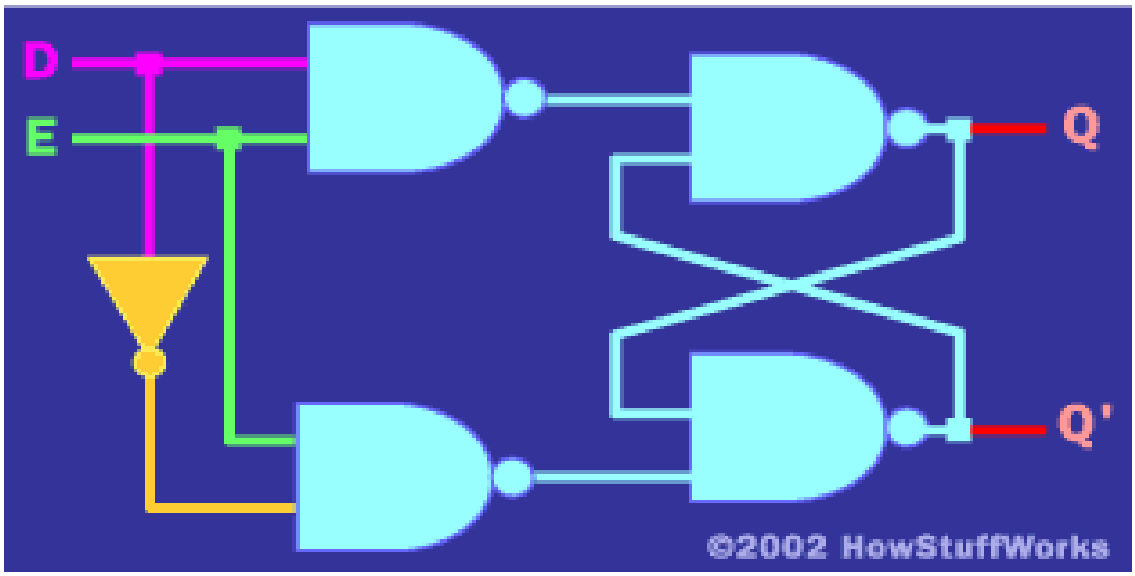

Figure 3. A SRAM memory cell consists of five to seven transistors.

Static RAM is a type of RAM that holds its data without external refresh, for as long as power is supplied to the circuit (figure 3). SRAMs are used for specific applications within the PC, where their strengths outweigh their weaknesses compared to DRAM:

Simplicity: SRAMs don't require external refresh circuitry or other work in order for them to keep their data intact.

Speed: SRAM is faster than DRAM.

In contrast, SRAMs have the following weaknesses:

Cost: SRAM is, byte for byte, several times more expensive than DRAM.

Size: SRAMs take up much more space than DRAMs (which is part of why the cost is higher).

These advantages and disadvantages taken together obviously show that performance-wise, SRAM is superior to DRAM, and we would use it exclusively if only we could do so economically. Unfortunately, 32 MB of SRAM would be prohibitively large and

costly, which is why DRAM is used for system memory. SRAMs are used instead for level 1 cache and level 2 cache memory, for which it is perfectly suited; cache memory needs to be very fast.

SRAM is manufactured in a way rather similar to how processors are: highly-integrated transistor patterns photo-etched into silicon. Each SRAM bit is comprised of between four and six transistors, which is why SRAM takes up much more space compared to DRAM, which uses only one (plus a capacitor). Because an SRAM chip is comprised of thousands or millions of identical cells, it is much easier to make than a CPU, which is a large die with a non-repetitive structure. This is one reason why RAM chips cost much less than processors do. The process is similar (but simplified somewhat) for making memory circuits.

3. Flash Memory

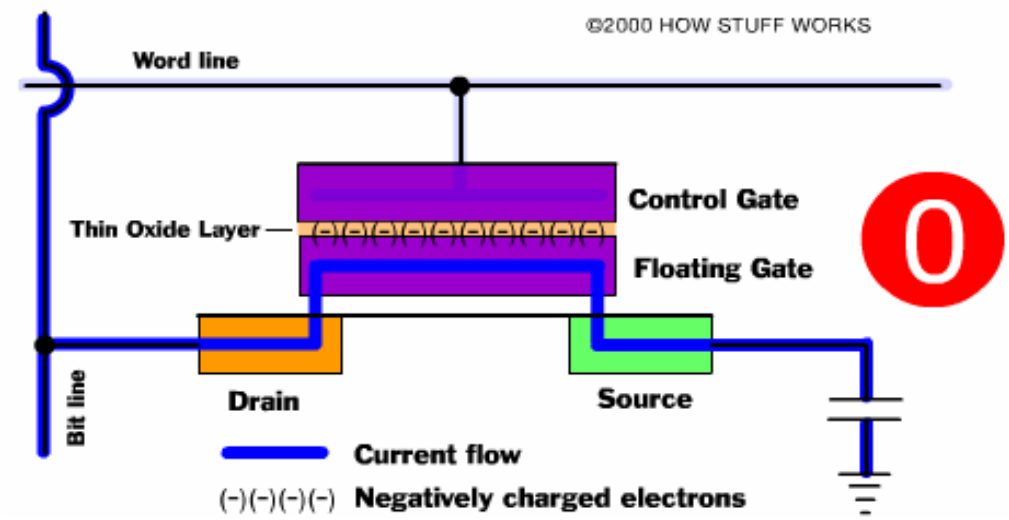

Figure 4. A flash memory cell has two transistors, sandwiching a thin oxide layer.

Flash memory has a grid of columns and rows with a cell that has two transistors at each intersection (figure 4). The two transistors are separated from each other by a thin oxide layer. One of the transistors is known as a floating gate, and the other one is the control gate. The floating gate's only link to the row, or wordline, is through the control gate. As long as this link is in place, the cell has a value of 1.

Tunneling is used to alter the placement of electrons in the floating gate. An electrical charge, usually 10 to 13 volts, is applied to the floating gate. The charge coming from the column, or bitline, enters the floating gate and drains to a ground.

This charge causes the floating-gate transistor to act like an electron gun. The excited electrons are pushed through and trapped on the other side of the thin oxide layer, giving it a negative charge. These negatively charged electrons act as a barrier between the control gate and the floating gate. A special device called a cell sensor monitors the level of the charge passing through the floating gate. If the flow through the gate is greater than 50 percent of the charge, it has a value of 1. When the charge passes through, the gate drops below the 50-percent threshold, the value changes to 0.

The electrons in the cells of a Flash-memory chip can be returned to normal ("1") by the application of an a higher-voltage electric field. Flash memory uses in-circuit wiring to apply the electric field either to the entire chip or to predetermined sections known as blocks. This erases the targeted area of the chip, which can then be rewritten.

The table below compares three important properties- access time, cell size and volatility, of the four different memories as mentioned above.

	MRAM	DRAM	SM	FM
Access time	15ns	60ns	10ns	10ms
Cell Size	0.6 μm^2	0.11μm^2	0.6μm^2	0.05μm^2
volatility	non-volatile	volatile	volatile	non-volatile

Current Induced Domain Wall Movement

One of the major problems MRAM faces is the memory cell size. At the current stage of development, the MRAM memory cell size is about five times larger than the size of the DRAM cell. The reason for the large cell size is that the Oersted fields generated by the current needed to switch the direction of the magnetic field in the magnet are too weak, so a large current density is required to make the switch. One possible way to solve this problem is by using current-induced domain-wall movement. An experiment performed by Stuart Parkin's group IBM showed the switching of the direction of the magnetic field inside a strip by passing polarized spin current through it[4].

Experimental setting

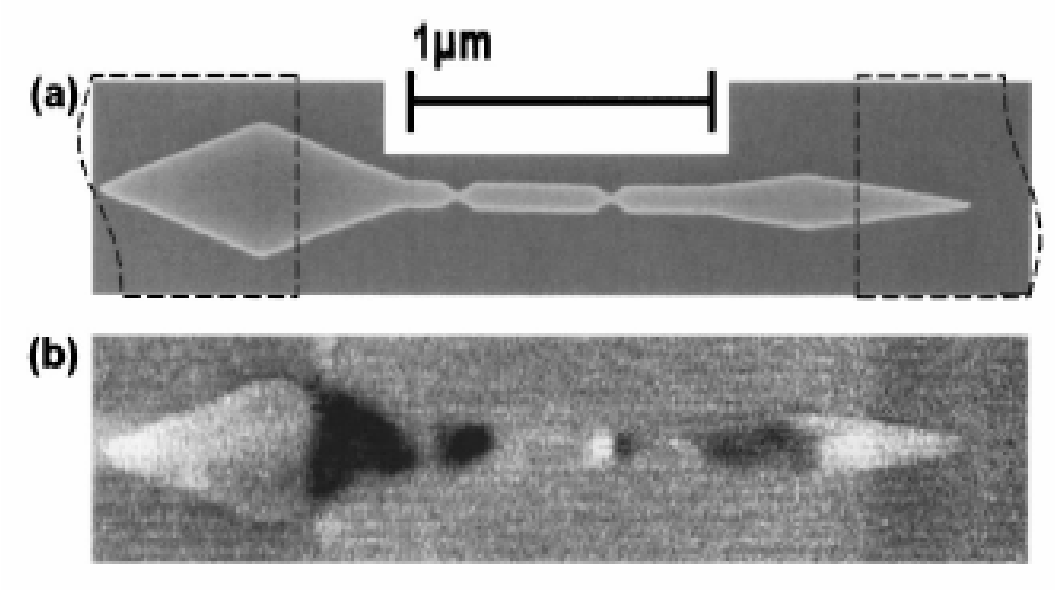

Figure 5. SEM (a) and MFM (b) image of the device. A domain wall is trapped in the left constriction[5].

Figure (5) shows a scanning electron microscope (SEM) image of the CoFe device. A tri-layer film of Ta 2.5 nm/CoFe 10 nm/Ta 2.5 nm was sputter deposited under ultrahigh vacuum conditions and then patterned using electron-beam lithography (EBL) and a subtractive ion mill process. The resulting CoFe shape is a 100 nm wide stripe connected to a large diamond shaped pad at the left end. The latter allows for the controllable injection of individual domain walls into the stripe; when the magnetic field is applied

along the stripe, the pad's magnetization reverses and a domain wall is injected into the stripe. Two constrictions ~50 nm wide in the stripe were used to trap the injected domain wall. Finally, two Au contacts, their overlap with the device is indicated by dashed lines in panel (a), were used to supply current and measure the resistance of the device. Note that the current flows primarily through the CoFe layer due to the much higher resistivity of the thin Ta layers.

The white area in figure (b) represents the magnetic field pointing to the left while the black area stands for the magnetic field pointing to the right. A domain wall is then formed between these two areas. When a spin polarized current passes through the system from left to right, the domain wall is pushed toward the right and finally disappears at the right end. Thus the direction of the magnetic field in the middle strip is flipped.

My work with Barbara Jones at IBM involves developing a model to explain the domain wall movement in the experiment above and finding out parameters related to the current that moves the domain wall.

Theoretical work to explain the domain wall movement

Consider a current propagating through a conducting ferromagnet (figure 6). Assume that there are two different kinds of electrons in the system, conducting electrons and localized electrons. Conducting electrons are free and interact only with local magnetization $\vec{M}$ determined by the total magnetic moment of the localized electrons. The motion of each individual conducting electron is governed by the Schrődinger equation with a term $J_H \vec{M}(r,t) \cdot \vec{\sigma}$ (1), where J_H is the value of the Hund's rule coupling or in general of the local exchange and $\vec{\sigma}$ is the spin of the conduction electron . Since spin-up electrons have lower energy, a nonzero average spin of the conduction electrons develops. An angular momentum density is then carried with the electric current so we have a flux of angular momentum. This leads to a nonzero average torque acting on the local magnetization which can deflect it from its original direction. The propagation of

current in a ferromagnet should be described by a coupled system of two equations: one for the motion of the conducting electrons and another for the magnetization of the localized electrons. We derive here the second equation, in the limit of small spacetime gradients, and present static and time dependent solutions for a 1D strip as shown in the experiment. The case of very large J_H is considered, meaning a complete polarization of the conduction electron spins in the ferromagnet. This case can often be realized in experiment. In layered structures, magnetic layers can be made of a material with a large band splitting, like a Heusler alloy, in which the spin opposite to the magnetization direction cannot propagate. In the Colossal Magneto Resistance materials, a large Hund's rule coupling is well known[6] and constitutes the basis for a double-exchange mechanism governing their magnetic ordering.

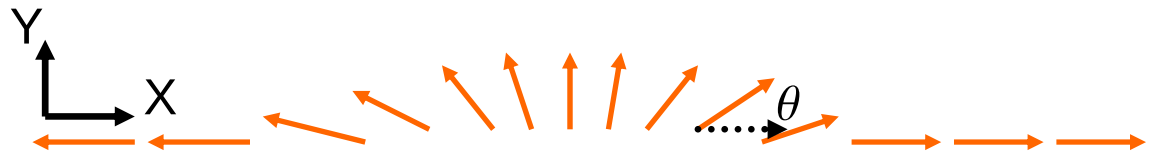

Figure 6. Spin-polarized current enters a half-infinite magnet from the left. Originally the magnetization is aligned along the easy-axis (-x). However if the incoming electrons are spin-polarized in a different direction, their interaction with the magnetization leads to a change of the direction of the local magnetization.

Schrödinger equation: the conduction electrons are considered to be non-interacting, so their wave-function satisfies,

$$i\hbar\frac{\partial\varphi}{\partial t}=-\frac{\hbar^2}{2m}\Delta\varphi+J_H\vec{M}(r,t)\cdot\vec{\sigma}\varphi\,, \qquad (2)$$

where $\varphi=\begin{pmatrix}\varphi_+\\ \varphi_-\end{pmatrix}$ and, $\vec{\sigma}$ is a Pauli Matrix. $\vec{M}=M\vec{n}$ with M defining magnitude of the local magnetic moment, and $\vec{n}$ being a unit vector in the direction of the local magnetic moment. And θ(r,t) is the angle between the local magnetic moment and the x axis.

We first, diagonalize the matrix $\vec{M}(r,t)\cdot\vec{\sigma}$ with a local spin rotation

$$U(r,t)=\begin{pmatrix}U_{++}(r,t) & U_{+-}(r,t)\\ U_{-+}(r,t) & U_{--}(r,t)\end{pmatrix}.$$

Define the rotated wave function ϕ via $\phi = U\varphi$, so that

$$\varphi = U^{-1}\phi, \tag{3}$$

then the spinor $\phi = \begin{pmatrix} \phi_+ \\ \phi_- \end{pmatrix}$ describes the conduction electron in a coordinate system with the z axis parallel to the local magnetization. Substituting (3) back to Eq.(2), while taking into accout $J_H \to \infty$, yields

$\phi_- \to 0$

$$ih\frac{\partial \phi_+}{\partial t} = -\frac{\hbar^2}{2m}\Delta\phi_+ + J_H M\phi_+ - \frac{\hbar^2}{2m}\sum_{j=\pm}(U_{+j}\nabla U_{j+})\nabla\phi_+ \tag{4}$$

Let $$A_i^{eff} = \frac{\hbar^2}{2m}\sum_{j=\pm}(U_{+j}\nabla_i U_{j+}) \tag{5}$$

The form of the equation (4) is the same as that of an electromagnetic interaction, and hence we can write by analogy

$$E_{\text{int}} = \mu_0 j_i A_i^{eff}, \tag{6}$$

with j the electric current and μ_0 the magnetic permeability.

The dynamics of the local magnetization is described by the Landau-Lifshitz equations[7] which are obtained from an energy functional. After adding Eq. (7) to the usual energy density of a ferromagnet with uniaxial anisotropy along the easy axis x.

$E = \int \left(J(\nabla \vec{M})^2 - K(x \cdot \vec{n})^2 \right) dV$, we obtain

$$E = \int \left(J(\nabla \vec{M})^2 - K(x \cdot \vec{n})^2 + \mu_0 j_i A_i^{eff} \right) dV, \tag{7}$$

where J is the exchange integral and K is the anisotropy constant.

The time dependence of the magnetization is now found from an equation of motion

$$\frac{\partial \vec{M}}{\partial t} = \frac{\mu_0 g|e|}{2m_e}\left[-\frac{\delta E}{\delta \vec{M}} \times \vec{M} \right] \tag{8}$$

Eq.(7) is the Landau-Lifshitz equation, which describes how the magnetic moment ($\vec{M}$) evolves in the effective field ($\vec{H}^{eff} = -\frac{\delta E}{\delta \vec{M}}$). In our case,

$$\vec{H}^{eff} = -\frac{\delta E}{\delta \vec{M}} = J\Delta\vec{M} + \frac{2K}{\vec{M}^2}(x\cdot\vec{M})x + \frac{\hbar}{2m}\frac{j_i}{e}[\nabla_i \vec{n}\times\vec{n}] \qquad (9)$$

There are three terms in the effective field. The first two are the exchange field and anisotropy field, respectively. The third term comes from the interaction between the conduction electrons and the local magnetic moments. Plugging Eq.(9) back into Eq.(8) finally yields

$$\frac{\partial\vec{M}}{\partial t} = \frac{\mu_0 g|e|}{2m_e}\left[J\Delta\vec{M} + \frac{2K}{M^2}(x\cdot\vec{M})x + \frac{\hbar}{2m}\frac{j_i}{e}[\nabla_i \vec{n}\times\vec{n}]\right]\times\vec{M} \qquad (10)$$

which is the starting point of our analysis.

In the stationary case, $\frac{\partial\vec{M}}{\partial t} = 0$ and j=0. All spatial derivatives vanish except in the x direction. In the experimental setting, the easy axis is also along the x direction. We derive an equation for n(r,t):

$$\tilde{J}\left[\frac{\partial^2\vec{n}}{\partial x^2}\times\vec{n}\right] = \left[\left(-\tilde{k}(x\cdot\vec{n})x\right)\times\vec{n}\right] \qquad (11)$$

with parameters

$$\tilde{J} = \frac{g\mu_B M}{\hbar}J, \tilde{K} = \frac{2g\mu_B}{\hbar M}K$$

In this case, the effective field, which has only exchange and anisotropy terms is anti-parallel to the local magnetic moment (figure 7).

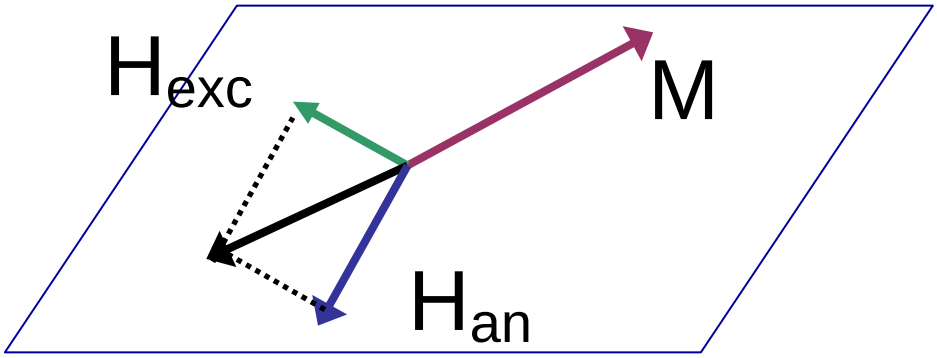

Figure 7. The vector sum of the exchange effective field (H_{exc}) and the anisotropy effective field (H_{ani}) is anti-parallel to the local magnetic moment (M).

The boundary condition at the right end of the sample requires that when x goes to infinity, $\vec{n}$ should align with the direction of the anisotropy, which is along the x axis. The solution to Eq.(11) is

$$\frac{\partial\theta}{\partial x} = \sqrt{\frac{\widetilde{K}}{\widetilde{J}}\left(1-\cos^2\theta\right)} \qquad (12)$$

With θ being the angle the magnetization makes with the x axis as shown in figure 6. The graph below shows the numerical results for the material PtMnSb (figure 8).

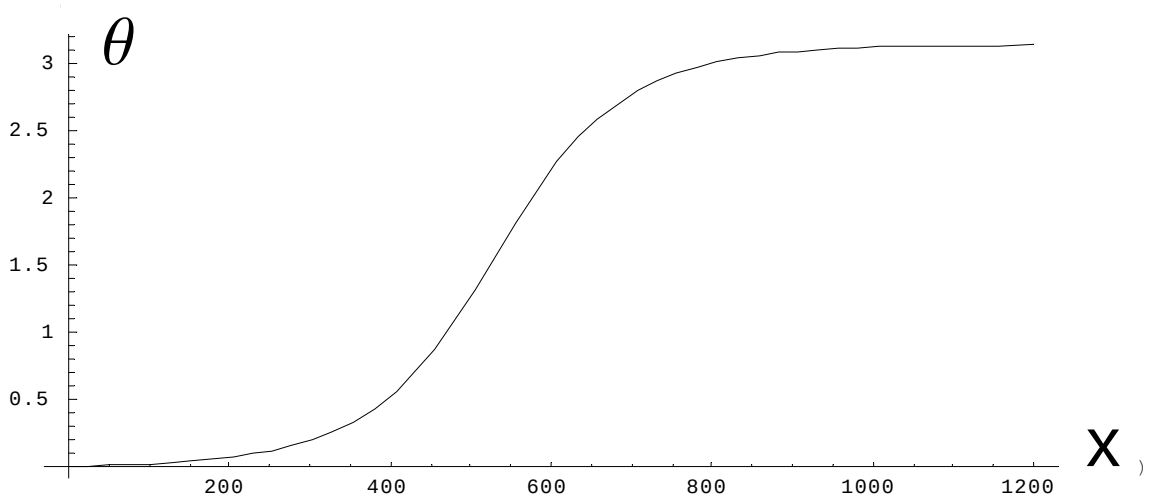

Figure 8. the numerical results shows a domain wall formed inside the strip. (and the parameters for PtMnSb are $K=1.19\times10^4\ J/m^3$, $\sqrt{\frac{\widetilde{J}}{\widetilde{k}}} = 2.77\times10^{-8}\,m$)

Next consider the time-dependent solution. $\frac{\partial\vec{M}}{\partial t}$ in Eq.(10) is not zero in this case. All the spatial derivatives also vanish except in the x direction (see figure 9). From Eq.(10), we derive

$$\frac{\partial\vec{n}}{\partial t} = \left[\left(\widetilde{J}\frac{\partial^2\vec{n}}{\partial x^2} + \widetilde{K}(x\cdot n)x\right)\times\vec{n}\right] - (Q\left(\frac{j_i}{e}\right)\frac{\partial\vec{n}}{\partial x}\times\vec{n})\times\vec{n}, \qquad (12)$$

where $Q = \frac{g\mu_B}{2M}$.

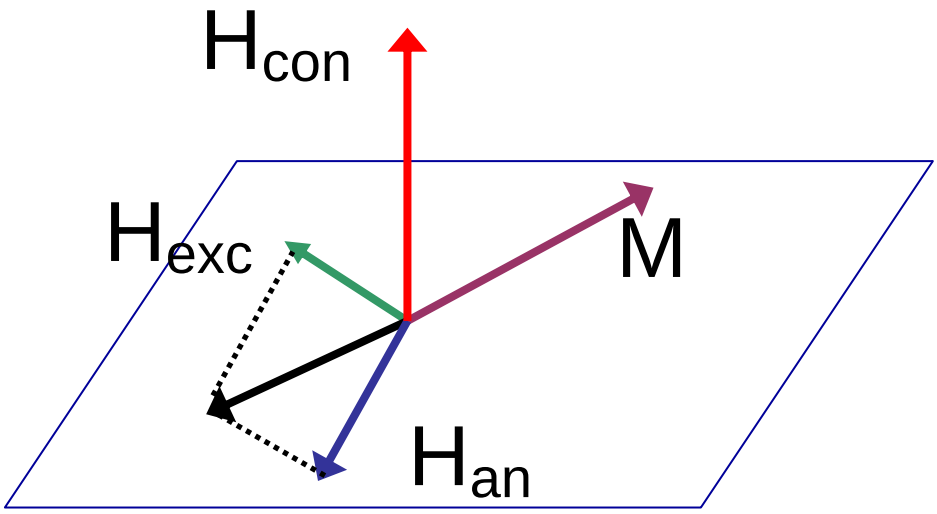

Figure 9. In the time dependent case, the vector sum of the exchange field (H_{exc}) and the anisotropy effective field (H_{ani}) is still anti-parallel to the local magnetic moment (M) as in the stationary case. The conducting electron effective field (H_{con}) is perpendicular to the plane.

The effective field here has three terms in the time dependent case-the exchange field, the anisotropy field and the conducting electron field (figure 10). The last term, the conduction electron field ($-Q\left(\frac{j_i}{e}\right)\frac{\partial n}{\partial x}\times n$) is perpendicular to the plane of the strip. This field try to align the neighboring local magnetic moments in the same direction. The field goes to zero when the local magnetic moments point in the same direction.

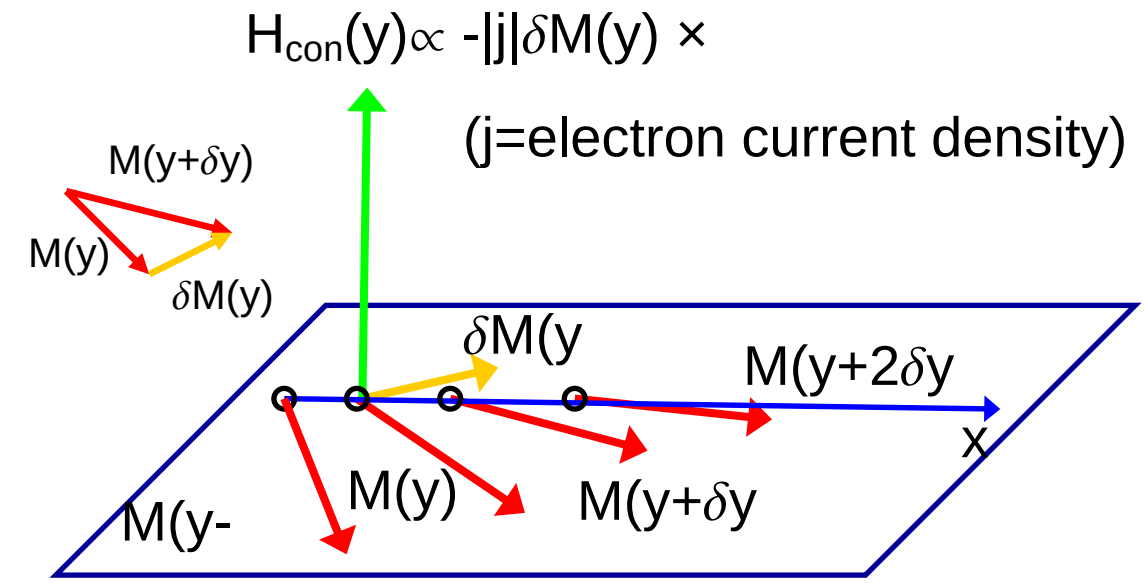

Figure 10. The conduction electron effective field (H_{con}) is proportional to the magnitude of the electron current density (j). The direction of the field is determined by the cross product of the vector difference between the neighboring local magnetic moments (δM(x)=M(x+δ)-M(x)) and M(x).

One of the time-dependent solutions is

$$n(r,t) = n_0[r + Q(\frac{j}{e})t] \qquad (13)$$

where n_0 is the solution in Eq.(12) for the static case. So this time-dependent solution describes the domain wall being pushed by the spin polarized current (figure 11).

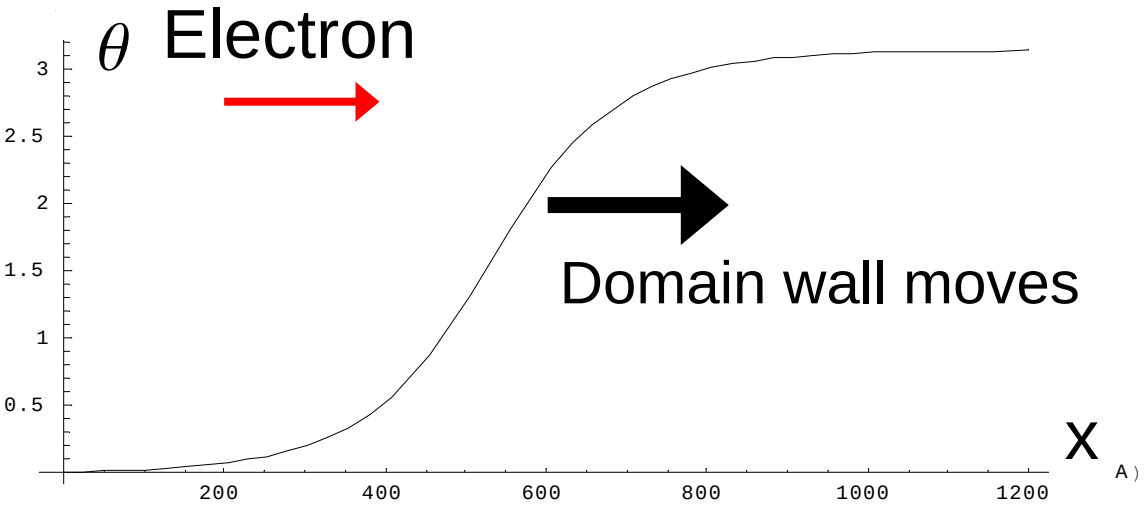

Figure 11. The numerical results of the time dependent case shows that the domain wall moves to the right as the polarized electrons are injected from the left to the strip.

The table below compares some numeric results for the domain wall size of specific materials with their corresponding experimental results.

Material	$LaCaMnO_3$ ($K=0.586\times10^4 J/m^3$, $\sqrt{\frac{\tilde{J}}{\tilde{k}}} = 2.77\times10^{-8} m$)	PtMnSb ($K=1.19\times10^4 J/m^3$, $\sqrt{\frac{\tilde{J}}{\tilde{k}}} = 4.50\times10^{-8} m$)
Experiment	>130nm	50-100nm
From the model	150 nm	80nm

Other works during the apprenticeship

For the first six months during my apprenticeship, under the guidance of Jim Freericks and Barbara Jones, I completed a publication entitled "Charge-density-wave order parameter of the Falicov-Kimball model in infinite dimensions"**,** (Phys. Rev. B. **68**,153102(2003)).

The Falicov-Kimball model[8] consists of two types of particles: itinerant conduction electrons and localized ions. The conduction electrons can hop between nearest-neighbor sites. The localized ions have an on site energy. There is a Coulomb interaction between the localized ions and the conduction electrons that sit at the same lattice site.

In the large-U limit (U is the Coulomb interaction between the localized ions and the conduction electrons that sit at the same lattice site), the Falicov-Kimball model maps onto an effective Ising model, with an order parameter described by a BCS-like mean-field theory in infinite dimensions. In the small-U limit, van Dongen and Vollhardt[9] showed that the order parameter assumes a strange non-BCS-like shape with a sharp reduction near T=Tc/2. In this publication, we investigated the crossover between these two regimes and quantitatively determined the order parameter for a variety of different values of U. We found the overall behavior of the order parameter as a function of temperature to be quite anomalous.

The evolution from the small-*U* to large-*U* limits is nonmonotonic. As *U* increases, the order parameter curve becomes steeper and steeper near *Tc* and flatter and flatter for small *T* until we reach the *U* corresponding to approximately the maximum of the *Tc* versus *U* curve. Once *U* increases past the maximum, then the order parameter curve slowly approaches the large-*U* limit.

We illustrate this behavior in the figure below for the Bethe lattice. When *U* lies below 0.2 [panel (a)], we can see the depression in the curve develop as we evolve toward the $U=0$ limit. For $0.2<U<2.6$, the curve moves to the right and becomes steeper [panel (b)]. Finally, for larger values of *U*, the curve moves to the left until it assumes the mean-field-theory form [panel (c)].

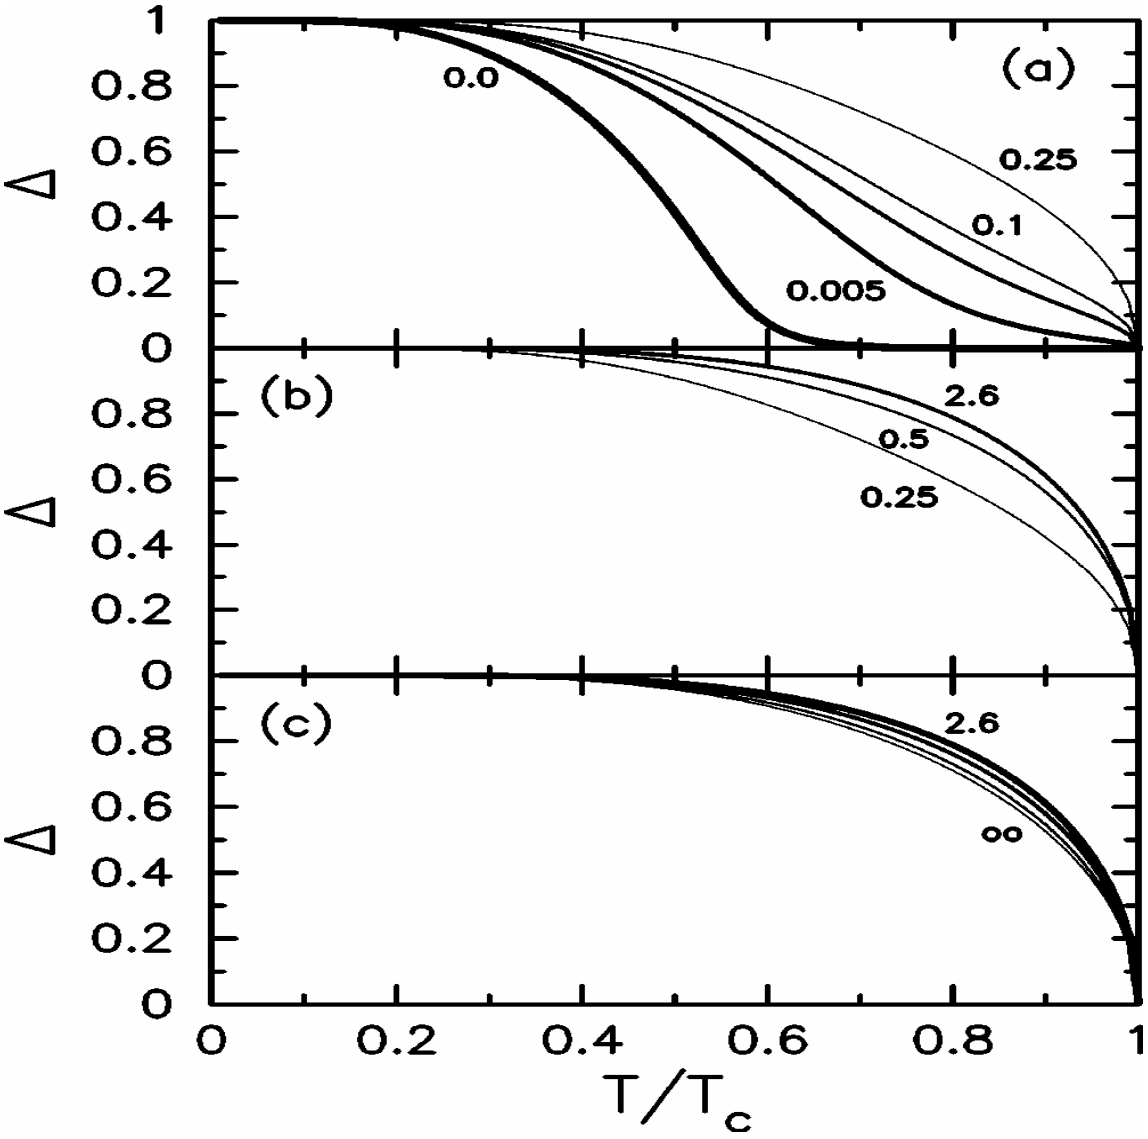

Figure 11.Order parameter for the CDW of the FK model at half filling on the Bethe lattice. Panel (a) shows the small-U results where a depression can be seen (U=0, 0.005, 0.03, 0.1,and 0.25), panel (b) shows the intermediate results (U=0.25, 0.5,and 2.6), where the order-parameter curve moves out to the right as U is increased, and panel (c) shows the large-U regime (U=2.6,3.0, 4.0, and $\propto$), where the curve shifts to the left and ultimately takes the mean-field-theory form.

I also spent several weeks on a parallel computing project. The purpose of doing that was to train myself in parallel computing. A parallel computing network combines the computing power of several computers together, speeding up the execution of certain programs by separating the task into several processes and running them in parallel. It is a

comparatively cheap way to achieve large computing power. There are several hardware architectures for a parallel network. Judging from the hardware available, I chose distributed memory architecture. The hardware I used for the network consisted of one HP laptop and two IBM pcs. The software installed for the parallel network included Redhat linux 9.0 (operating system), SSH (communication package) and MPICH[10](message passing library). I successfully built my own parallel computing network after connecting the network, installing the entire software package and configuring the parameters. I wrote several programs to perform numerical integration of simple functions like sin and cos functions using my parallel computing network.

---

[1] Hoover's corporate report www.hoover.com

[2] David Talbot, TR. **104**, 1 (2001)

[3] Part of the information about the memories is from www.howstuffworks.com

[4,5] M. Tsoi, R. E. Fontana, and S. S. P. Parkin, Appl. Physics. Lett. **83**, 13 (2003)

[6] A. R. Bishop and H. Roder, *Current Opinion in Solid State and Materials Science* (Current Chemistry Ltd., London, 2000)

[7] Eugene Chudnovsky and Javier Tejada, *Introduction to Magnetism* (Rinton Pr Inc., New York, 1996)

[8] L. M. Falicov and J. C. Kimball, Phys. Rev. Lett. **22**, 997 (1969)

[9] P. G. J. Van Dongen, Phys. Rev. B **45**, 2267 (1992). P. G. J. van Dongen and D. Vollhardt, Phys. Rev. Lett. **65**, 1663 (1990).

[10] www-unix.mcs.anl.gov/mpi/mpich/