RESEARCH REVIEW & ADVISORY COUNCIL MEETING

2018

CMRR

Center for Memory and Recording Research

THURSDAY & FRIDAY NOVEMBER 1-2, 2018

CMRR @ UC SAN DIEGO

cmrr.ucsd.edu

RESEARCH REVIEW PROGRAM

Thursday, November 1

8:30 AM	Continental breakfast at CMRR	
8:55 AM	Welcome and Introduction	
9:00 AM	**Tribology, Mechanics and Medical Device Technology**	Professor Frank Talke
	Investigation of Voltage Biasing and Relative Humidity on Tribochemistry of the Head-Disk Interface	Tan Trinh
	Life Time Studies of HAMR Technology	Tan Trinh
	Flutter and Non-repeatable Runout of Disks in Air, Vacuum, and Helium	Christoph Schade, Jason Zeng
	A Wireless Handheld Pressure Measurement System for in vivo Monitoring of Intraocular Pressure in Rabbits	Alex Phan
10:00 AM	10 MINUTE BREAK	
	Anchoring Methods for Implantable IOP sensors	Phuong Truong
	Esophageal Deflection Device	Karcher Morris
	Internet-Enabled Patient Point-of-Care Devices for Monitoring Eye Conditions	Buu Truong, Nick Williams
	Electricidal Urinary Catheter	Oren Gotlib
	Designing a Detachable Bronchoscope	Matthew Kohanfars
11:30 AM	**Magnetic Films and Nanostructures**	Professor Eric E. Fullerton
	Recovery of antiferromagnetic order in Cr following photoexcitation	Sheena Patel
	Electronic Transport Properties of Au/Si Multilayers	Yuxuan Xiao
	Intermediate Temperature Thin Film Solid Oxide Fuel Cell (SOFC)	Haowen Ren
	New Opportunities for Ultrafast Spintronics	Professor Eric E. Fullerton
12:35 PM	LUNCH - CMRR LOBBY	

RESEARCH REVIEW PROGRAM

Thursday, November 1

1:35 PM	**Materials and Electronic Devices**	Kenji Nomura
	Oxide semiconductor for next generation electronics	
2:00 PM	**Thermal Energy Transport**	Professor Renkun Chen
	Probing Thermal Transport mediated by Surface Phonon Polariton	Sunmi Shin
	Managing Ultrahigh Heat Flux using Thin Film Boiling	Qingyang Wang
2:30 PM	**Unconventional Computing with Memory**	Massimiliano (Max) Di Ventra
	MemComputing: a CMRR Success Story	
3:00 PM	10 MINUTE BREAK	
3:10 PM	**Micromagnetic Modeling & Recording Physics**	Professor Vitaliy Lomakin
	Computing resonant modes and excitation states in micromagnetic systems with a finite-element based frequency domain solver	Zhuonan Lin
	Fast Landau-Lifshitz-Gilbert Equation- Valet Fert Equations Finite Element Method Solver	Xueyang Wang
	Micromagnetic modeling of non-uniformities in magnetic tunnel junctions for MRAM Devices	Iana Volvach
	Simulation of Domain Wall Displacement Using the Landau Lifshitz Lambda Model	Marco Menarini
	Macromagnetic Simulation of THz signals in Antiferromagnetic FeRh	Marco Menarini
5:10 PM	**Poster Session**	
5:30 PM	**Reception**	
6:00 PM	**Advisory Council Meeting**	

RESEARCH REVIEW PROGRAM

Friday, November 2

8:30 AM	Continental breakfast at CMRR	
8:55 AM	Welcome and Introduction	
9:00 AM	**Signal Processing & Coding**	Professor Paul Siegel
	Flash Memory from a Broadcast-Channel Perspective	Yonglong Li
	Shaping Codes for Costly Channels with Memory: Theory and Applications	Yi Liu
	ECC Layouts for Memory Systems with Variation	Wei Wu
	Experimental Soft Bit Read Applications	Yi Liu
	NAND Flash Channel Characterization	Navya Sree Prem
10:30 AM	10 MINUTE BREAK	
10:40 AM	**Non-Volatile Solid State Memory**	Professor Steven Swanson
	Finding and Fixing Performance Pathologies in Persistent Memory Software Stacks	Juno Kim
11:10 AM	**Special Session: Quantum Materials for Energy Efficient Neuromorphic Computing (Q-MEEN-C), EFRC** Oleg Shpyrko Professor, Department of Physics	
11:35 AM	**Special Session : Seminar** Jishen Zhao Assistant Professor, Computer Science and Engineering Department	
12:15 PM	LUNCH - CMRR LOBBY	

Investigation of Voltage Biasing and Relative Humidity on Tribochemistry of the Head-Disk Interface

Presenter: ***Tan D. Trinh***, Ph.D. student, CMRR

Researchers: ***Tan D. Trinh***, Ph.D. student, CMRR

Christoph Schade, Visiting graduate student, CMRR

Collaborator: ***Michael Johnson***, Ph.D., Seagate Technology

Advisor: ***Frank E. Talke***, Professor, CMRR

In current hard disk drives, the spacing between the recording head and the rotating disk is less than 1 nm. At such a small spacing, intermittent contacts between the head and the media greatly affect the tribological properties of the head-disk interface, such as wear and lubricant degradation. Recently, applied bias voltage has been shown to improve the tribology of the head-disk interface [1-3]. Rajauria et al. [2] and Matthes et al. [3] showed that applying a negative bias voltage to the slider with respect to the disk improves the wear performance of the slider surface compared to the situation when a positive bias voltage is applied. Tani et al. [4] and Tan et al. [5] studied the effect of bias voltage on the lubricant transfer across the head-disk interface. They showed that lubricant transfer is a strong function of the bias voltage between head and disk.

In this study, we investigate the effect of voltage biasing and environmental relative humidity on the tribological properties of the head disk interface. A commercially available load/unload tester (VENA tester) located inside an environmental chamber was used for testing. A source-meter (Keithley, 2400B) was used to apply a bias voltage to the spindle with the slider body being grounded. A heater element embedded inside the slider was used to control the flying height of the slider. An acoustic emission sensor, glued to the actuator arm was used to detect contacts between the head and the disk. The contact potential across the head-disk interface was determined by measuring the power-to-contact (PtC) as a function of applied bias voltage. "Accelerated" burnish testing was performed to study the effect of bias voltage and environmental relative humidity on the change of the power-to-contact, i.e., to study wear of the head-disk interface. Pitting corrosion was observed in the read and write area of the slider surface and is a strong function of relative humidity level and the polarity of the applied bias voltage across the head-disk interface.

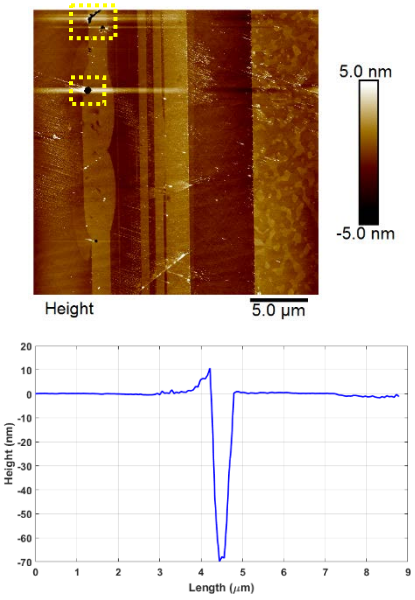

References

[1] Knigge, B., Marchon, B., "Negative biasing a slider with respect to a disk to reduce slider wear and provide burnish rate control", US Patent 8,139,309 B2, 2012.

[2] Rajauria, S., Schreck, E., Marchon, B., "Voltage assisted asymmetric nanoscale wear on ultra-smooth diamond like carbon thin films at high sliding speeds", Nature Scientific Reports, 2016.

[3] Matthes, L.M., Spada, F. E., Ovcharenko, A., Knigge, B.E., Talke, F.E., "Effect of Head-Disk Interface Biasing and Relative Humidity on Wear of Thermal Flying Height Control Sliders", Tribology Letters, 2017.

[4] H. Tani, S. Koganezawa, N. Tagawa, "Reduction in lubricant pickup by bias voltage between slider and disk surfaces", Microsystem Technologies, 2016

[5] B. K. Tan, B. Liu, Y. Ma, M. Zhang, S. F. Ling, "Effect of Electrostatic Force on Slider-Lubricant Interaction", IEEE Transactions on Magnetics, 2007

Life-time Studies of Heat-Assisted Magnetic Recording Technology

Presenter: ***Tan D. Trinh***, Ph.D. student, CMRR

Researcher: ***Tan D. Trinh***, Ph.D. student, CMRR

Collaborators: ***Sukumar Rajauria***, Ph.D., Western Digital

Erhard Schrek, Ph.D., Western Digital

Advisor: ***Frank E. Talke***, Professor, CMRR

Heat-assisted magnetic recording (HAMR) is a leading technology contender to extend the areal density of hard disk drives beyond 1Tb/in^2 [1]. In heat-assisted magnetic recording, light from a laser diode mounted on the top of the slider is delivered through a wave guide to a plasmonic near-field transducer (NFT) that is used to locally heat the media surface. A nanoscale area of the media surface is heated up to near the Curie temperature, thereby reducing the coercivity of the magnetic layer to near zero. At this low a coercivity, data can be written on the disk using current design magnetic transducer. After writing, the temperature decays quickly, "freezing" the magnetization on the previously written bit. The high temperature needed for recording poses many challenges on the reliability of the head-disk interface, including concerns related to heat transfer, head-media contacts and lubricant degradation [2].

In this study, a fully integrated system of HAMR heads and HAMR media was used to study life-time of the head-disk interface using a commercially available read-write tester (Guzik Technical Enterprises, USA). A burst writing scheme [3] was used to estimate writer pole tip protrusion (WPTP) and near-field transducer protrusion (NFTP) as shown in the figure on the right. A clearance of 2nm was maintained during testing using an embedded heater element located inside the slider. In this study, an increase of write track width (MWW) was used to accelerate the failure of the head-disk interface. Failure of the head-disk interface was based on a 1dB change of signal to noise ratio (ΔSNR) measurements, i.e., a significant change in the signal amplitude. Our study shows that HAMR head life-time is a strong function of the writing track width (MWW) and the optical power of the laser diode.

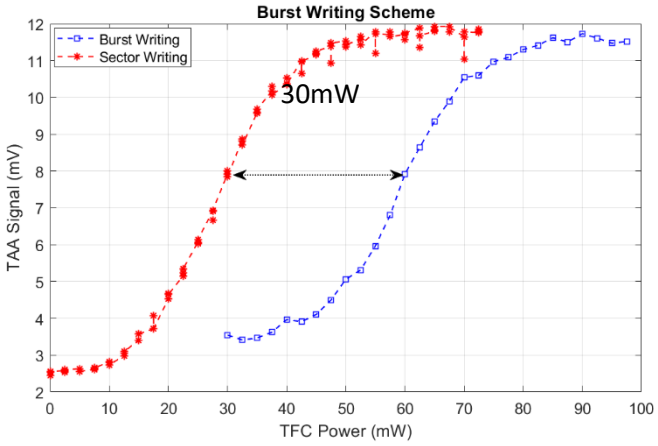

References

[1] Kryder M. et. al., "Heat Assisted Magnetic Recording", Proceedings of the IEEE, 2008

[2] Zhang J. et. al., "Lubrication for Heat-Assisted Magnetic Recording Media", IEEE Transactions on Magnetics, 2006

[3] Xiong S. et. al., "Setting Write Spacing in Heat Assisted Magnetic Recording", IEEE Transactions on Magnetics, 2018

Flutter and Non-repeatable Runout of Disks in Air, Vacuum, and Helium

Presenter: ***Christoph Schade***, Visiting student, TU Dresden, CMRR

Researchers: ***Tan D. Trinh***, Ph.D. student, CMRR
Zijian Zeng, Ph.D. student, CMRR
Zhengqiang Tang, Visiting Professor, CMRR

Collaborators: ***Karim Kaddeche***, L2 Drive, Irvine, CA
John Wang, L2 Drive, Irvine, CA

Advisor: ***Frank E. Talke***, Professor, CMRR

In current hard disk drives, the spacing between the read/write elements and the media surface is approximately 1 nm. To increase the data density above 1Tb/in^2 [1], the spacing must be reduced further, a very difficult task. One approach recently proposed by Kaddeche [2] is to place the hard disk inside a vacuum environment and servo the head in the vertical direction without the presence of an air bearing. A vacuum environment reduces media runout, and poses a potentially improved situation for servoing the head above the disk in the absence of an air bearing. In this study, we investigate the repeatable and non-repeatable runout of a rotating disk in vacuum and helium conditions using a laser Doppler vibrometer and a capacitance probe.

A 3.5-inch (97mm) single-platter hard drive without an actuator arm was used for testing. The drive was placed in a chamber, in which the environmental condition were controlled using a vacuum pump and a pressure valve. The rotational speed of the media is controlled using a DC motor controller. The disk vibrations were measured using a laser Doppler vibrometer (Polytec, Germany) and capacitance probe (Microsense, USA). The runout of the media surface were investigated as a function of the rotational speed of the disk, the radius and the vacuum level. Our results showed that as the pressure decreases, the repeatable runout of the media surface increased while the non-repeatable runout decreased. In addition, the repeatable runout is shown to be a strong function of the rotational speed of the disk and its radius.

[1] Marchon B., Pitchford T., Hsia Y-T., Gangopadhyay S., "The Head-Disk Interface Roadmap to an Areal Density of 4Tbit/in^2", Advances in Tribology, 2013
[2] Karim K., "Active Control of a Read/Write Head for reduced head-media spacing", US Patent 9666229B1

A Wireless Handheld Pressure Measurement System for *in vivo* Monitoring of Intraocular Pressure in Rabbits

Presenter: ***Alex Phan***, Postdoctoral Researcher, CMRR, MAE

Researchers: ***Alex Phan***, Postdoctoral Researcher, CMRR, MAE
Phuong Truong, Graduate Student, CMRR, MAE
Kerrianne Stewart, Research Assistant, CMRR, MAE
Benjamin Suen, Graduate Student, CMRR, MAE

Collaborators: **Andrew Camp M.D.**, Assistant Professor, Shiley Eye Institute
Robert Weinreb M.D., Professor, Shiley Eye Institute

Advisor: ***Frank E. Talke***, Professor, CMRR, MAE

Glaucoma is a neurodegenerative disease that affects 65 million people worldwide, causing progressive optic nerve damage over time and irreversible blindness. Elevated intraocular pressure in the eye has been considered one of the leading modifiable risk factors of the disease. The objective of this study is to evaluate an optical pressure measurement system for frequent and direct measurement of intraocular pressure *in vivo*.

The IOP measurement system is comprised of a submillimeter (< 0.3 mm^3) size implantable pressure sensor and an optical readout unit. Once the IOP sensor is implanted, the optical reader can optically measure pressure by generating and capturing interference patterns from the sensor. Two sensors were tested *in vitro* to characterize the behavior in a pressure range from 0 mmHg to 100 mmHg. The sensitivity and accuracy of the sensor *in vitro* was 30 nm/mmHg and ±0.2 mmHg, respectively. The sensors then were integrated into an intraocular lens and implanted during cataract surgery in two rabbits. During the ten weeks study, the two New Zealand white rabbits responded well to the implant with no signs of inflammation, infection, or protein deposition. We successfully acquired pressure readings from minimally restrained animals using the portable DSLR handheld reader (Fig. 1). IOP measurements obtained from the implanted sensor agreed well with the measurements taken from a tonometer (a device currently used in clinics to measure IOP) in both long-term study and transient study (Fig. 2).

The results of this study show promising capability of measuring direct and accurate intraocular pressure conveniently and remotely using an optical sensor. Frequent pressure measurements will enable physicians to make well informed diagnosis and treatment plans and allows patients to better manage their condition.

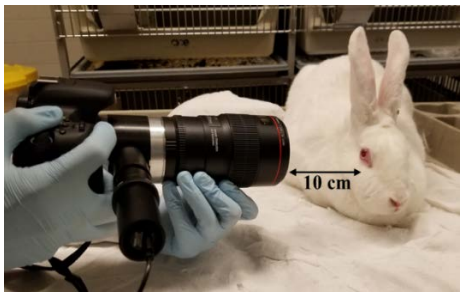

Figure 1. Intraocular pressure measurement taken using the portable DSLR reader

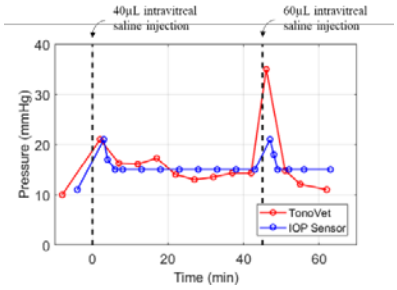

Figure 2. Results of *in vivo* transient intraocular pressure study

Anchoring Methods for Implantable IOP sensors

Presenter: ***Phuong Truong***, PhD Student, Mechanical Engineering

Researchers:
Phuong Truong, PhD Student, Mechanical Engineering
Alex Phan, Postdoctoral Researcher, Mechanical Engineering
Kerrianne Stewart, Research Assistant, Mechanical Engineering
Buu Truong, PhD Student, Bioengineering
Dr. Weinreb, Professor of Ophthalmology, Shiley Eye Institute
Dr. Andrew Camp, Ophthalmologist, Shiley Eye Institute

Advisor: ***Frank Talke,*** Professor, Mechanical Engineering

Glaucoma is considered the second leading cause of blindness worldwide with an estimated 112 million cases by 2040 [1]. To date, intraocular pressure (IOP) is considered as a primary risk factor for the progression of the disease and is generally measured in clinic using a Goldmann applanation tonometer.

In recent years, many implantable intraocular pressure sensors have been developed to overcome challenges related to inaccuracies of standard tonometry techniques by directly measuring pressure. However, sensor anchoring mechanisms and methods have not been thoroughly evaluated to consider standard surgical techniques and optimal implant location.

In this study, we present various anchoring methods for small implantable pressure sensors that do not disrupt standard ocular surgery procedures. Namely, attaching the sensors to other implants such as a capsular tension ring (Fig. 1), a rigid artificial cornea (Fig. 2), an intraocular lens (Fig. 3), and an iris claw (Fig. 4). Using the intraocular pressure sensor previously developed by Phan et al., each implant device was integrated with the pressure sensor and tested ex vivo using enucleated rabbit eyes. Initial proof of concept results has demonstrated promising results with reliable measurement readings.

Ultimately, integrating the pressure sensor will enable patients to monitor their eye pressure and improve glaucoma management, patient compliance, and help preserve the vision of the patients.

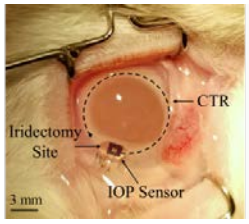

Fig. 1. Sensor integrated with capsular tension ring implanted using an iridectomy procedure.

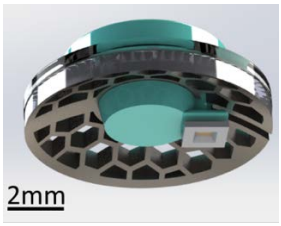

Fig. 2. Proposed modified Kpro design with sensor embedded.

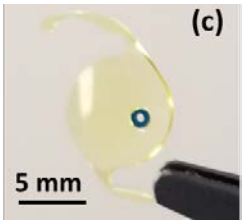

Fig. 3. An intraocular lens with sensor embedded into the lens.

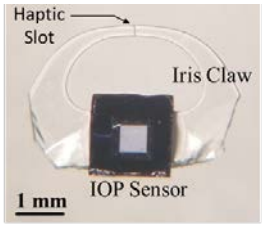

Fig. 4. A modified phakic lens with iris claw end removed and integrated with intraocular pressure sensor.

Esophageal Deflection Device

Presenter: **Karcher Morris**, Ph.D. Student, CMRR & MAE

Researchers: **Andreas Rosenkranz**, Post-Doctoral Researcher, CMRR

Scott Garner, Graduate Researcher, CMRR & MAE

Dora Trieu, Undergraduate Researcher, CMRR & MAE

Collaborators: **Thomas Savides**, Medical Doctor, UCSD Health System

Gregory Feld, Medical Doctor, UCSD Health System

Advisor: **Frank Talke**, Professor, CMRR & MAE

Less than two years ago, our research team was presented with a significant medical problem of thermal esophageal damage occurring during cardiac ablation procedures. Since then, we have been working on a device to deflect the esophagus away from the heart during the sensitive procedure. The novel device, illustrated in figure 1, is simply comprised of an overtube with a curved end and an insertion piece. The assembled device can be inserted into the esophagus and the insertion piece can be pulled to activate the shape memory of the overtube to move the esophagus to safety (figure 2).

Prototypes have been developed utilizing advanced 3D printing capabilities assisted by a detailed understanding of material properties and finite element simulation. Current manufacturing methods for the device take into consideration mass production with medical grade materials as well as incorporated temperature measurement features. Ex vivo studies have been performed to explore the devices ability to deflect the esophagus as well as measure heat transfer through heart and esophagus tissue. Future work will be discussed with an emphasis on plans to perform clinical trials at UCSD's Jacobs Medical Center.

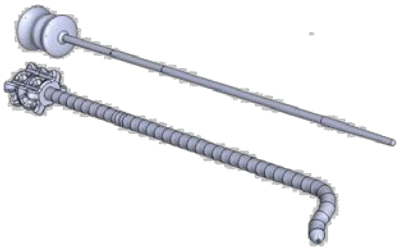

Fig. 1: CAD model of the Esophageal Deflection Device (EDD)

Fig. 2: 3D Print of EDD illustrating deflection

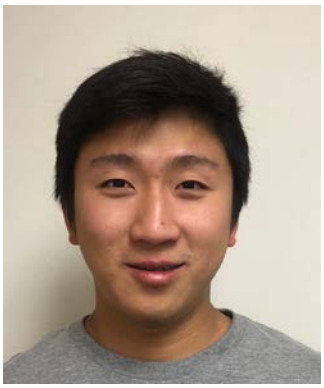

Internet-Enabled Patient Point-of-Care Devices for Monitoring Eye Conditions

Presenters: ***Buu Truong***, Graduate Student, MAE
Nick Williams, Undergraduate Student, MAE

Researchers: ***Buu Truong***, Graduate Student, MAE
Nick Williams, Undergraduate Student, MAE
Alex Phan, Postdoctoral Researcher, MAE
Phuong Truong, Graduate Student, MAE
Dr. Kang Zhang, Ophthalmologist, Shiley Eye Institute

Advisor: ***Frank Talke,*** Professor, MAE

Internet-enabled patient point-of-care devices are transforming the way doctors interact with their patients, especially in the area known as teleophthalmology. Technologies in teleophthalmology allow patients to remotely connect with physicians to monitor their eye conditions. By enabling frequent monitorization and timely therapeutic intervention, these devices save time and money for both the patient and doctor.

In this project, we develop two tele-ophthalmic (internet-enabled ophthalmic) devices that can be used to monitor the ocular conditions: (1) a portable smartphone-based handheld reader to acquire intraocular pressure readings and (2) a funduscope phone attachment for patient point of care imaging of the retina. When used in conjunction with an intraocular pressure sensor, the portable IOP reader allows the user to self image their eye and take accurate pressure readings. These frequent measurements provide the physicians with more information about the patient's condition, which is pivotal for disease management. As a scientific tool, the reader will allow researchers to gain insight on glaucoma and accurately measure pressure. Additionally, the funduscope leverages the ubiquity of smartphone and its camera to allow retinal examination outside of clinical offices.

Preliminary prototyping of both devices have demonstrated very promising results. They are compact, easy to use, and affordable. Fig. 1 illustrates the design of a compact IOP reader. Fig. 2 shows the fringe patterns that the reader detects, and then uses software to determine IOP. Fig. 3 shows the model of a funduscope, and Fig. 4 is the retinal image obtained using the device.

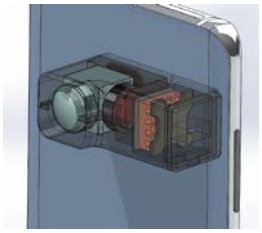

Fig. 1. Compact IOP handheld smartphone reader.

Fig. 2. Interference fringe captured from smartphone reader. Used to calculate IOP.

Fig. 3. Portable funduscope

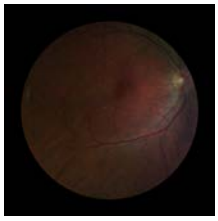

Fig. 4. Retinal image from smartphone funduscope.

Electricidal Urinary Catheter

Presenter: ***Oren Gotlib***, Masters Student, CMRR & MATS

Researchers: ***Fred Spada***, Associate Research Scientist, CMRR

Adriane Minori, Ph.D. Candidate, MAE

Karcher Morris, Ph.D. Candidate, CMRR & MAE

Collaborators: ***Madhu Alagiri,*** M.D, UCSD Health System

Katy Patras, Postdoctoral Fellow, UCSD School of Med.

Mike Tolley, Associate Professor, MAE

Advisor: ***Frank Talke***, Professor, CMRR & MAE

Catheter acquired urinary infection or CAUTI is the most prevalent type of healthcare related infection, accounting for more than 30% of acute care hospital infections and associated with 13,000 deaths each year. Great efforts have been made to fight CAUTI by modifying catheters with bactericidal coatings, microscopic patterns and even electrical current. Last summer, we assembled an interdisciplinary team to understand the electricidal effect in order to translate it into an infection resistant catheter. Our *in vitro* studies have shown the electricidal effect in a commercial wound dressing works to inhibit biofilm with patterned deposits of known antimicrobials, zinc and silver oxide (figure 1). These studies confirm that the design produces electric potentials in solution that generate both bactericidal electric fields and hydrogen peroxide. Our team has begun to synthesize solutions of these galvanic active materials and we are making efforts to incorporate them onto catheter surfaces to fight CAUTI (figure 2 & 3).

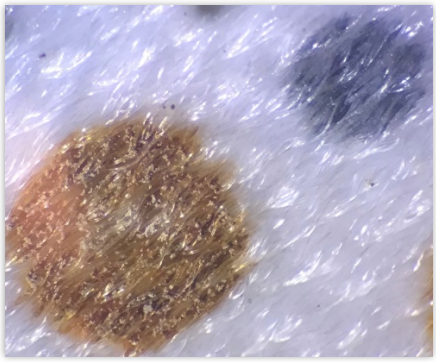

Figure 1: Bioelectric wound dressing (20x Magnification)

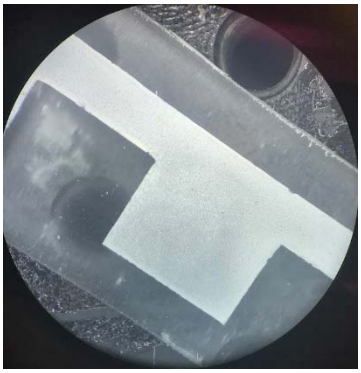

Figure 2: Screen-printed bioelectric film on PDMS (5x Magnification)

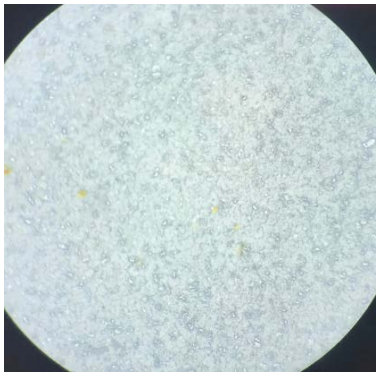

Figure 3: Screen-printed bioelectric film on PDMS (20x Magnification)

Designing a Detachable Bronchoscope

Presenter: ***Matthew Kohanfars***, Masters student, CMRR & MAE

Researchers: ***Karcher Morris***, Ph.D. Candidate, CMRR & MAE

Benjamin Suen, Masters student, CMRR & MAE

Collaborators: ***Dr. Jaspreet Somal***, Medical Doctor, UCSD Health System

Advisor: ***Dr. Frank Talke***, Professor, CMRR & MAE

Intubation is a common and extremely sensitive procedure performed by an anesthesiologist. During the process of intubation, a bronchoscope is used to place an endotracheal tube into the trachea to maintain an open airway for an anesthetized patient. A detachable bronchoscope would be highly desirable for patients and physicians with challenging and unforeseen events during procedures. Our research team has taken steps to design and develop a practical, reliable, and economical medical device. In the search for the optimal design, prototypes have been designed using electromagnets, flexible hooks, clevis pins, and spring loaded mechanisms. 3-D models were designed, tested, and animated using SolidWorks to search for defects and potential enhancements to the overall design (Figure 1). From knowledge gained through iterations of machined and 3D printed prototypes, a spring-loaded detachment mechanism is, as of yet, the most promising design.

Fig. 1: CAD model of a bronchoscope

Recovery of antiferromagnetic order in Cr following photoexcitation

Presenter: ***Sheena K.K. Patel***, PhD Student, Physics

Researchers: ***Sheena K.K. Patel, Eric Fullerton, Andrej Singer, Oleg Shpyrko***

Advisor: ***Eric Fullerton***, Professor, ECE

The emergence of order in systems such as antiferromagnets, superconductors, and density-wave systems results in intriguing phenomena such as spin density waves and symmetry-lowered ordered states. Exploration through ultrafast excitation can provide insight into the mechanisms driving these phenomena. In Cr, we study the charge-density-wave (CDW) order of Cr which appears as the second harmonic of the spin-density-wave (SDW) ordering and is measured in x-ray diffraction as satellite peaks about the (002) peak, *e.g.* at (0, 0, 2±2δ), where 2δ is the momentum transfer of the CDW (see Fig. 1). As a function of temperature and film thickness, the satellite peak will align to different Laue fringes on either side of the Bragg peak, indicating a change in the number of periods of the CDW between the substrate interface and the surface of the sample. In our film, the CDW is aligned to the 8th or 7th fringes at different temperatures, indicating 8.5 or 7.5 CDW periods respectively. We conducted an X-ray pump probe experiment at the X-ray free electron laser LCLS, where a 28 nm Cr film was excited by optical (800 nm, 40 fs) pulse and probed with an x-ray photon energy of 8.9 keV with a temporal resolution of 50 fs. Following strong excitation of the system at different temperatures, we are able to measure the change in the magnetic order including the discrete change in CDW and SDW period and the suppression in order when we pump above the Neel temperature. When pumping through the hysteretic change in CDW/SDW period, the magnetic order recovers to a metastable state that persists for more than 400 ns after pumping.

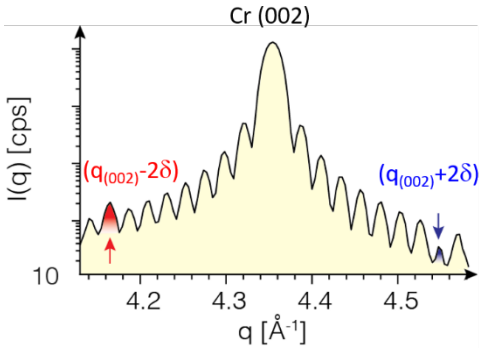

Figure 1: X-ray diffraction signal from a 28 nm Cr film, showing satellite peaks due to the charge density wave ordering.

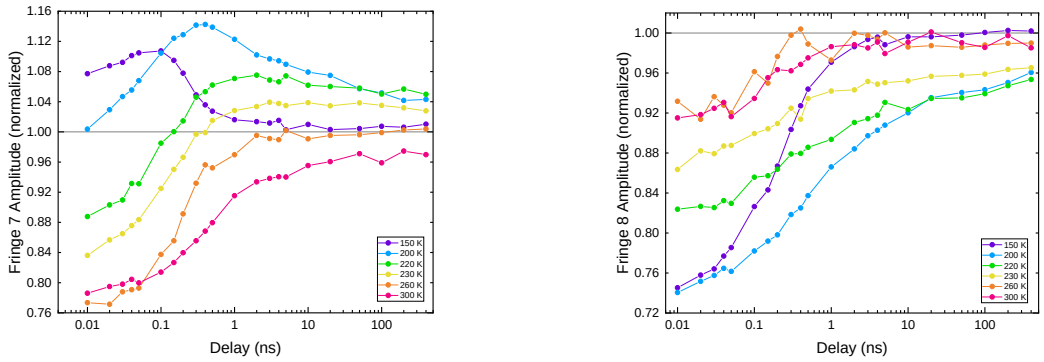

Figure 2: Fringes 7 and 8 amplitude following photoexcitation normalized to the pre-pump amplitude. At temperatures below 200 K, the CDW satellite peaks are aligned to fringe 8, indicating 8.5 CDW periods. Between 200 K and 240 K, the CDW satellite disappears from fringe 8 and appears on fringe 7. The Neel temperature is near 290 K.

Electronic Transport Properties of Au/Si Multilayers

Presenter: ***Yuxuan Xiao,*** PhD Student, MatSci

Researchers: ***Yuxuan Xiao***, ***Mohammed Salah El Hadri, Eric Fullerton***

Advisor: ***Eric Fullerton,*** Profesor, CMRR and ECE

The interplay of electron and spin transport in ultrathin metal films gives rise to new phenomena that may impact emerging technologies. Both anisotropic magnetoresistance (AMR) [1] and Hall effect [2] have been used to probe the quantum size effects in ultrathin metal films. Cherradi et al. first investigated the 2D effects in ultrathin Au/Si multilayers via the AMR [1].

We will report on the transport properties of Au (t_{Au})/Si (t_{Si}) multilayers with 0.5 nm $\leq t_{Au} \leq$ 4 nm and 0.5 nm $\leq t_{Si} \leq$ 8 nm. The multilayer nature and intermixing were studied using low angle x-ray reflectivity measurements. The ordinary Hall resistance shows pronounced variations of the Hall coefficient (R_H) with the Au thickness suggest the presence of quantum size effects in the ultrathin individual Au layers. (see Fig.1)

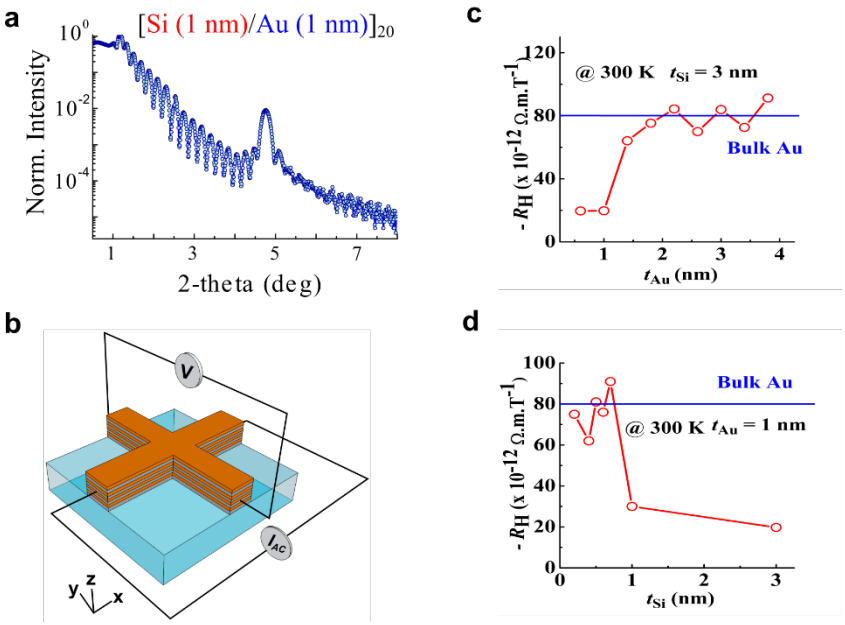

Fig. 1: (a) Low angle θ-2θ x-ray reflectivity from [Si (1 nm)/Au (1 nm)]$_{20}$ multilayers. (b) Schematic of the ordinary Hall measurement. (c) R_H as function of the thickness of the individual Au layers t_{Au} for [Si (3 nm)/Au (t_{Au})]$_7$ multilayers (d) R_H as function of the thickness of the individual Si layers t_{Si} for [Si (t_{Si})/Au (1 nm)] multilayers.

References:

[1] N. Cherradi, A. Audouard, G. Marchal *et al*. *Phys. Rev. B* **39**, 7424 (1989).

[2] M. Jalochowski, M. Hoffman, and E. Bauer, *Phys. Rev. Lett.* **76**, 4227 (1996).

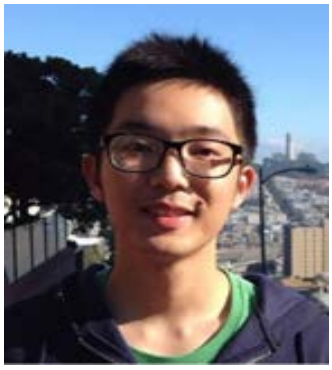

Intermediate Temperature Thin Film Solid Oxide Fuel Cell (SOFC)

Presenter: ***Haowen Ren,*** Ph.D. Student

Researchers: ***Yoon-Ho Lee, Erik Wu, Tuyen Tran, Eric E. Fullerton, Nguyen Minh, Shirley Y. Meng***

Advisor: ***Eric Fullerton***, Professor, CMRR, ECE

The crisis in petrol and also the wary of global warming, the green energy technologies, like battery and fuel cell etc., has emerged as an ideal alternative to fossil fuels. Solid Oxide Fuel Cell (SOFC) is one of the most promising energy technology in this kind. Contrary to the commonly known battery which is a type of energy storage device, fuel cell is an energy conversion device that transfer chemical energy directly into electricity with high efficiency.

The conventional SOFCs are typically fabricated by compressing, sintering and annealing, which usually requires processing temperature over 1000 ^{o}C for days. Instead of using conventional fabrication methods, our research utilize magnetron sputtering technology to fabricate the thin film SOFC cell in a single chamber at room temperature. The advantages of thin film SOFC fabricated by sputtering are obvious: less amount of materials consumed, fast approach process, room temperature deposition and single-chamber procedure. Since the reducing of electrolyte thickness, the oxygen ionic resistance greatly suppressed which lowering the operation temperature from 900~1000 ^{o}C to 600~700 ^{o}C. At this intermediate temperature range, stainless steel interconnect is robust enough to support the cell stacks. At the same time, the nano-column structure of anode and cathode largely increased the triple phase boundaries where electrochemistry reactions take place.

Porous anode is Ni/YZr co-sputtering, as Ni is catalyst and YZr (oxidized to YSZ at air) is supporting and ionic conducting materials. Dense thin YSZ are deposited as electrolyte and thin interlayer of GDC for functional layer. Porous LSCF-YSZ is co-sputtering as cathode for oxygen reduction. Pt is deposited as current collector. The peak performance about 2.1 W/cm^2 is the best sputtered thin film SOFC performance ever in existed published references at 650 ^{o}C. Ethanol has been tested as an alternative fuel source, which shows 250 mW/cm^2 at 600 ^{o}C. Methane is the targeted fuel source in our study.

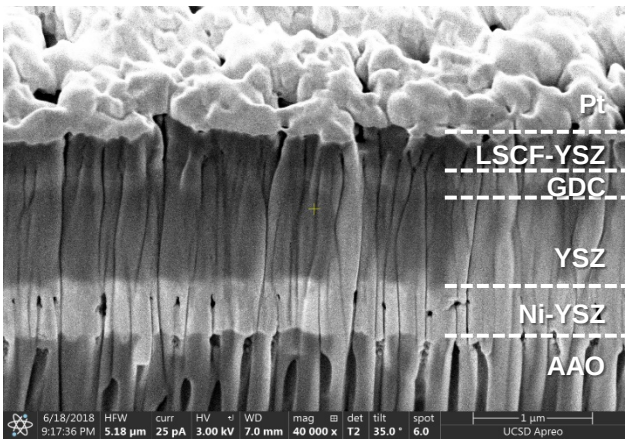

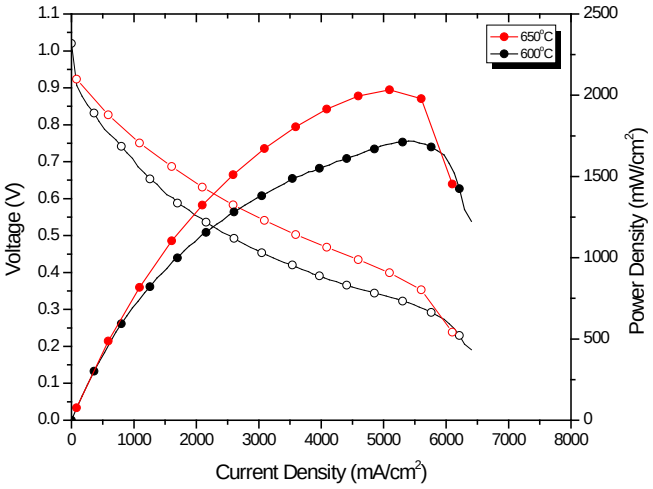

Oxide semiconductor for next generation electronics

Presenter: ***Kenji Nomura,*** Assistant Professor, Electrical and Computer Engineering

Amorphous metal oxide semiconductors represented by a-In-Ga-Zn-O (a-IGZO) are widely recognized as a promising material to develop next generation flexible, transparent and low-cost electronic devices because of its superior material properties such as optical transparency in visible region, reasonably high electron mobility (>>10cm^2(Vs)$^{-1}$) and wide compatibility of processing such as low temperature dc sputtering and low-cost solution process including printable/rollable technology. Therefore, high-performance thin film transistor (TFT) using a-IGZO as n-channel have been developed for next generation display applications and mass production of AMFPDs such as high-resolution OLEDs and low-power consumption displays have started with a-IGZO-TFT backplane.

Next challenge is to develop high performance p-type oxide and to realize p-channel TFT, flexible complimentary circuit, and several *pn*-junction devices such as photovoltaic, light-emitting devices and so on. However, oxide semiconductors with excellent electrical property are all n-type due to the unique nature of electronic structure composed of n*s* orbital (n is quantum number) of cation for conduction band minimum (CBM) and 2*p* orbital of oxygen with strong directivity for valence band maximum (VBM). Therefore, hole transport in many oxides is easily degraded due to the VBM nature. The absence of high-performance p-type oxide is largest drawback in oxide device. Therefore, it is imperative to develop high performance p-type oxide material and the related electric devices such as p-channel TFTs.

Here I will present the material design for p-type oxide semiconductor and development of p-channel oxide-TFTs. Firstly, I will introduce amorphous oxide semiconductor and oxide TFT technology. Then I will present high-performance p-channel oxide-TFT and the oxide-based complimentary inverter applications.

Probing thermal transport mediated by surface phonon polariton

Presenter: ***Sunmi Shin,*** Ph.D. Student, Materials Science

Advisor: ***Renkun Chen***, Assoc Professor

Coherent thermal emission deviates from the thermal emission predicted for a blackbody by the classical Planck's law. Although the existence of polaritonic radiation and the associated relative enhancement of the thermal emission in far-field have been previously explored through modelling and optical characterizations, an approach to achieve and directly measure dominant coherent thermal emission has not materialised. We exploited the large disparity in the two length scales pertaining to thermal emission in polar dielectrics: the skin depth and resonance wavelength of surface phonon polaritons (SPhPs). We then designed anisotropic nanoribbons made of polar dielectrics (e.g., SiO_2) to enable independent control of the incoherent and coherent behaviours. We observed over 8.5-fold enhancement in the emissivity of the nanoribbons compared with the thin-film limit. Importantly, this enhancement is attributed to the coherent resonant effect of the SPhPs and, hence, was found to be more pronounced at lower temperature. A novel thermometry platform was devised to extract, for the first time, the thermal emissivity from such nanoscale dielectric thermal emitters despite the perceived experimental challenge given the exceedingly low emitting power (0.1–10 nW). The demonstrated coherent heat source provides new insight into the realisation of spatial and spectral distribution control for thermal emission.

Managing ultrahigh heat flux using thin film boiling

Presenter: ***Qingyang Wang***, PhD Student, MAE

Researchers: ***Qingyang Wang, Renkun Chen***

Advisor: ***Renkun Chen,*** Assoc Professor, MAE

Phase change heat transfer is fundamentally important for thermal energy conversion and management, such as in electronics with power density over 1 kW/cm^2. The critical heat flux (CHF) of phase change heat transfer in theory is greater than 1 kW/cm^2 on a planar surface, but its experimental realization has remained elusive. Moreover, phase change heat transfer is usually not actively tunable which limits its application in thermal regulation. Here, we utilized nanoporous membranes to realize a new "*thin film boiling*" regime that resulted in an unprecedentedly high CHF of over 1.2 kW/cm^2 for water boiling on a planar surface. The thin film configuration reduces the conductive thermal resistance, leads to high frequency bubble departure, and provides separate liquid-vapor pathways, therefore significantly enhances the heat transfer. This new "thin film boiling" concept provides a new approach to achieve ultrahigh heat flux in phase change heat transfer and will benefit the application in thermal management of high power devices. Moreover, we demonstrated the active and reversible modulation of the thermal conductance of thin film boiling of ethanol. The conductance was widely tunable for over 100 times, which paves the way for constructing active thermal switches and regulators to dynamically control heat transfer for a variety of applications.

MemComputing: a CMRR success story

Presenter: ***Massimiliano Di Ventra***, Professor, Physics

Memcomputing [1, 2, 3, 4] is a novel computing paradigm that employs memory to both store and process information on the same physical location. The research behind this approach has been primarily funded by CMRR and has led to the founding of a start-up, MemComputing, Inc. (http://memcpu.com/), located in San Diego and already delivering solutions for some important players in industry. In this talk, I will discuss the power of memcomputing by showing results on a wide variety of combinatorial optimization problems [5, 6, 7] where simulations of these machines already demonstrate substantial advantages compared to traditional algorithms. I will also discuss the prospects of fabricating these machines in hardware with available technology, hence offering a realistic path to real-time computing.

[1] M. Di Ventra and Y.V. Pershin, *Computing: the Parallel Approach*, **Nature Physics** 9, 200 (2013).

[2] M. Di Ventra and Y.V. Pershin, *Just add memory*, **Scientific American** 312, 56 (2015).

[3] F. L. Traversa and M. Di Ventra, *Universal Memcomputing Machines*, **IEEE Transactions on Neural Networks and Learning Systems** 26, 2702 (2015).

[4] M. Di Ventra and F.L. Traversa, *Memcomputing: leveraging memory and physics to compute efficiently*, **J. Appl. Phys.** 123, 180901 (2018).

[5] F. L. Traversa and M. Di Ventra, *Polynomial-time solution of prime factorization and NP-complete problems with digital memcomputing machines*, **Chaos: An Interdisciplinary Journal of Nonlinear Science** 27, 023107 (2017).

[6] F. L. Traversa, P. Cicotti, F. Sheldon, and M. Di Ventra, *Evidence of an exponential speed-up in the solution of hard optimization problems*, **Complexity** 2018, 7982851 (2018).

[7] M. Di Ventra, F. L. Traversa and I.V. Ovchinnikov, *Topological field theory and computing with instantons*, **Annalen der Physik** 529,1700123 (2017).

Computing resonant modes and excitation states in micromagnetic systems with a finite-element based frequency domain solver

Presenter: ***Zhuonan Lin***, Ph.D. Student, Materials Science and Engineering

Advisor: ***Dr. Vitaliy Lomakin***, Professor, Electrical and Computer Engineering Department

Solving and implementing excitation states and resonant mode serve as a powerful tool applicable to various micromagnetic systems. Linearized time domain Landau-Lifshitz-Gilbert (LLG) equation can be used to study situations with weak excitations but has limitations when the solution at a particular frequency is of interest. We present our recently developmed finite element method (FEM) based linear solver in frequency domain. The solver can compute excitation states with an iterative procedure, and it is accelerated by a preconditioner. This method can be used for modeling linear phenomena, such as spin waves. We also implement an FEM-based eigen-value solver for predicting normal oscillation frequencies and modes in complex micromagnetic systems. The developed codes are useful for applications, such as calculating noise in magnetic sensors, magnetic memory systems, and spin torque oscillators. We carried out numerical experiments to show the validity of simulation results and to study various systems. The solver is integrated with the general-purpose FastMag simulator to make it efficient by using efficient methods to compute the effective fields and accelerate by Graphics Processing Unit (GPU) and multi-core CPU computing.

Fast Landau–Lifshitz–Gilbert Equation – Valet Fert Equations Finite Element Method Solver

Presenter: ***Xueyang Wang***, PhD Student, ECE

Researcher: ***Xueyang Wang***, PhD Student, ECE

Collaborator: ***Alexander Goncharov***, Researcher, Western Digital

Advisor: ***Vitaliy Lomakin***, Professor, ECE

Spin transfer torque (STT) is an effect providing the ability to modify the orientation of magnetization by spin polarized current. Spin polarized current can be generated by passing an unpolarized electric current through a magnetic layer. This current when passing through another layer transfers its angular momentum. The angular momentum transfer can change the orientation of the second magnetic layer or can excite oscillations. This effect is used in applications, such as STT-based Magnetic Random Access Memories (STT-MRAM) and Microwave Assisted Magnetic Recording (MAMR).

Here, we show a solver that couples the Landau–Lifshitz–Gilbert equation with the Valet-Fert theory for computing the spin accumulation in complex systems. The solver is based on Finite Element Method and this allows modeling highly complex magnetic structures. We present details of the formulation and implementation. The Valet-Fert solver is accelerated on both the algorithmic level, via a predictor, and the hardware level, via GPU and multi-core computing. Simulation results for modeling a MAMR head are presented.

Micromagnetic modeling of non-uniformities in magnetic tunnel junctions for MRAM Devices

Presenter: ***Iana Volvach***, PhD Student, CMRR, MSE

Advisor: ***Vitaliy Lomakin***, Professor, CMRR, ECE

Magnetic devices, such as magnetic random access memories (MRAM) experimentally have an arbitrary shapes with edge and surface roughness. Accurate models are needed to enable the simulation of novel magnetic devices with the roughness and arbitrary shapes. We developed the LLG simulator, which self-consistently models the dynamics of the magnetic device and present results of using this simulator for accurate modeling of MRAM switching to understand the effect of edge/surface roughness. The simulator uses FastMag micromagnetic solver [1] for considering complex dynamics in MRAM, accounting for magnetoresistance, non-uniform magnetization, finite temperature, and complex material and geometry composition. In the presented approach a magnetic device is represented as a for layers stack: free layer, reference layer and untiferromagneticaly coupled SAF1 and SAF2. SAF2 bottom thickness is tuned to minimize average field in free layer. It was noticed that the switching dynamics has noticeable effects on non-uniform switching dynamics, thermal stability, critical current density and figure of merit when the MRAM stack has the edge and surface roughness. Roughness makes stability 10-20 k_bT higher. Fig. 1. shows micromagnetic model results of critical current density dependence of MRAM size with and without edge roughness.

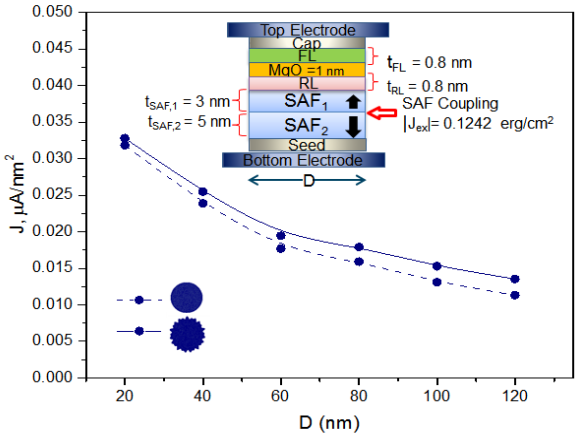

Fig. 1. a) Critical current density dependence of MRAM size element. The MRAM free layer has perpendicular surface anisotropy energy density Ks = 1 erg/cm2, Ms = 1432 emu/cm3, alpha = 0.008.

References:

1. R. Chang, S. Li, M. Lubarda, B. Livshitz, and V. Lomakin, “FastMag: Fast micromagnetic simulator for complex magnetic structures,” Journal of Applied Physics, vol. 109, p. 07D358, 2011.

Simulation of Domain Wall Displacement Using the Landau Lifshitz Lambda Model

Presenter: ***Marco Menarini***, PhD Student, CMRR, MSE

Advisor: ***Vitaliy Lomakin***, Professor, CMRR, ECE

Recent experiments have shown helicity dependent domain wall displacement (HD-DWD) when the domain between two states is excited by a series of low fluency sub-picosecond laser pulses[1,2]. To simulate this process, we developed a model coupling the Landau-Gilbert-Bloch (LLB) model with a three-level lambda system, which we refer to as the Landau-Lifshitz-Lambda model (LLL). The laser excitation drives the system magnetization to an excited state. This effect is followed by a quick decay into the reversed magnetization ground state. The optical relaxation is followed by a slower thermal decay via the classic LLB magnetization dynamics. The combination of this optical excitation and the thermal demagnetization leads to an asymptotical value for the displacement. We present simulations of HD-DWD in a continuous FePtL10 thin film. We compare our results with the ones obtained using MCD and the experimental results (Fig.1).

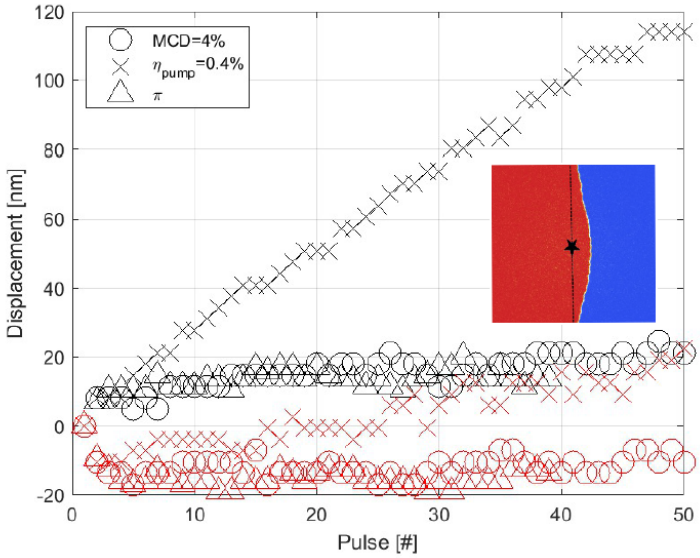

Fig. 1. HD-DWD due to linear polarized light (triangles), MCD (circles), and optical pumping (crosses) vs number of pulses respect to the center of the laser.

References:

[1] Y. Quessab, R. Medapalli, M. El Hadri, M. Hehn, G. Malinowski, E. Fullerton, and S. Mangin, arXiv preprint arXiv:1709.07645 (2017).

[2] R. Medapalli et al., arXiv:1607.02505 (2016).

Macromagnetic Simulation of THz signals in Antiferromagnetic FeRh

Presenter: ***Marco Menarini***, PhD Student, CMRR, MSE

Advisor: ***Vitaliy Lomakin***, Professor, CMRR, ECE

The use of antiferromagnets has been proposed as possible alternative to ferromagnetic material for producing THz signals using STT-induced damping compensation [1]. In this work, we studied the laser induced THz-signals in FeRh/Pt bilayer material using a micromagnetic model for Ferrimagnetic media [2] when subjected to sub-picosecond laser pulses. We choose the FeRh for its metamagnetic transition from antiferromagnetic (AFM) to ferromagnetic (FM) ordering above a transition temperature Tm. The micromagnetic exchange between the two sublattice of Fe has been obtained from atomistic simulation and experiment as a function of the temperature [3]. Our simulations show that transition between the FM-AFM transition of the exchange component generates the THz signal observed experimentally (Fig 1).

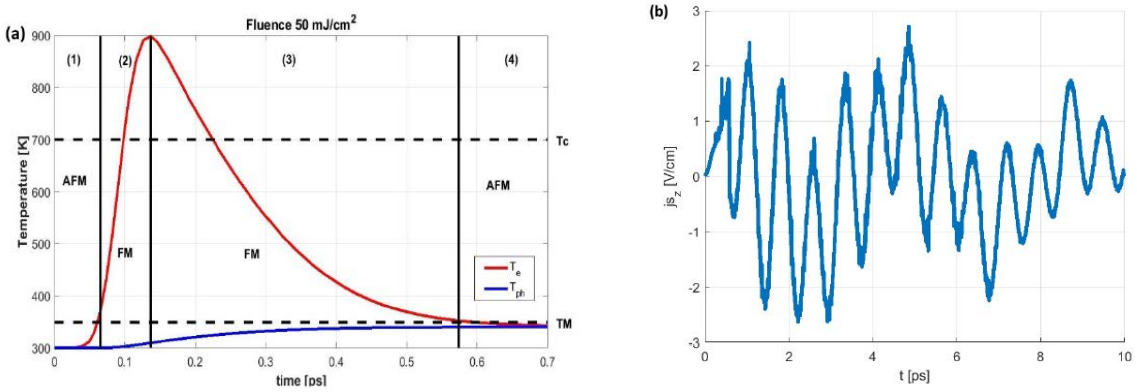

Fig. 1. (a) Shows the temperature profile of the material during the laser pulse. The two sublayers transition from an AFM coupling to a FM one during the heating process. When the temperature drops below the transition point TM, the THz signal is generated, as shown in (b).

References:

[1] R. Khymyn, I. Lisenkov, V. Tiberkevich, B. A. Ivanov, and A. Slavin, Scientific Reports 7, 43705 (2017).

[2] U. Atxitia, P. Nieves, and O. Chubykalo-Fesenko, Physical Review B 86, 104414 (2012).

[3] J. Barker and R. W. Chantrell, Physical Review B 92, 094402 (2015).

Flash Memory from a Broadcast-Channel Perspective

Presenter: ***Yonglong Li,*** Postdoc, CMRR, UCSD

Researcher: ***Yonglong Li,*** Postdoc, CMRR, UCSD

Collaborators: ***Pengfei Huang***, Staff Engineer, Western Digital Corp.

Advisor: ***Paul H. Siegel***, Professor, CMRR, UCSD

In this work we study the multi-level cell (MLC) NAND flash memory from the broadcast channel perspective. We first propose two broadcast channels to model MLC NAND flash memory and then numerically compute upper and lower bounds on the sum rate, which translates to the coding rate of MLC NAND flash memory.

Shaping Codes for Costly Channels with Memory: Theory and Applications

Presenter: ***Yi Liu***, Ph.D. Student, ECE

Researcher: ***Yi Liu***, Ph.D. Student, ECE

Collaborators: ***Yonglong Li***, Postdoctoral Researcher, CMRR, ***Pengfei Huang***, Graduate Student, ECE

Advisor: ***Paul H. Siegel,*** Professor, ECE

In our previous work [1], we studied shaping codes for flash memory channels with independent cost model. In this talk, we consider costly channel models that can incorporate inter-cell interference (ICI) [2], in which programming a flash cell to a high voltage level will also damage its neighbors that are programmed to a low voltage level. One example of an SLC flash memory channel with ICI is shown graphically in Fig. 1, where the vertices are overlapping pairs of two channel symbols and the edge labels represent the cost of a transition.

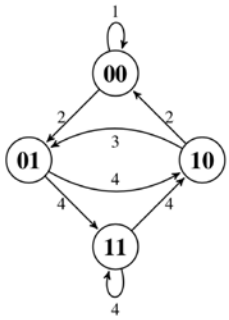

Fig. 1 SLC Flash Memory Channel: Costly Channel Model

In this talk we analyze and design shaping codes for the general costly channel with memory l. We first analyze shaping codes for a fixed code rate, which we call *type*-I *shaping codes*. We develop a theoretical bound on the trade-off between the rate and the average cost of a type-I shaping code. We then study shaping codes that minimize average cost per source symbol (total cost). We refer to this class of shaping codes as *type-II shaping codes*. We then define a modified costly channel where edge cost was determined by the optimal symbol occurrence and prove the equivalence theorem showing that both type-I and type-II shaping codes can be realized using a type-II shaping code for this modified channel. As an extension of Varn coding [3] for a memoryless costly channel, we present an asymptotically optimal type-II shaping code on this modified channel for a uniform i.i.d. source. A consequence of the analysis is a separation theorem, which states that optimal shaping can be achieved by a concatenation of lossless compression and optimal shaping for a uniform i.i.d. source. This provides an alternative architecture for implementing shaping codes.

[1] Y. Liu, P. Huang and Paul. H. Siegel, "Performance of Data Shaping Codes for Flash Memory", 2016 CMRR Fall Research Review, La Jolla, CA, Oct. 13-14, 2016.

[2] V. Taranalli, H. Uchikawa and P. H. Siegel, "Error Analysis and Inter-Cell Interference Mitigation in Multi-Level Cell Flash Memories," in *Proc. IEEE Int. Conf. Commun. (ICC)*, London, Jun. 2015, pp. 271-276.

[3] B. Varn, "Optimal variable length Codes (arbitrary symbol cost and equal code word probability)," *Inform. Contr.*, vol. 19, pp. 289-301, 1971.

ECC Layouts for Memory Systems with Variation

Presenter: ***Wei Wu***, Ph. D. Student, ECE, UCSD

Researcher: ***Wei Wu***, Ph. D. Student, ECE, UCSD

Advisor: ***Paul H. Siegel***, Professor, ECE, UCSD

We consider error correction code (ECC) design for memory systems with variation. The variation is an inevitable property of real systems because of the difference among devices. It is theoretically expressed through a channel model with different qualities. To deal with the variation, we first propose a probabilistic channel model, which is different from the standard algebraic model. Then, we describe the design of an ECC that can tolerate one device failure. By separately storing correlated segments of ECC codewords into different devices, we can apply a single ECC in layout designed for the system with variation. The layout is more efficient to implement in practice, with a comparable error-rate performance, compared to previous designs.

Experimental Soft Bit Read Applications

Presenter: ***Yi Liu***, Ph.D. Student, ECE

Researcher: ***Yi Liu***, Ph.D. Student, ECE

Advisor: ***Paul H. Siegel,*** Professor, CMRR

In this talk, we report on several experimental soft bit read applications. We first introduce the program/cycling experiment that we use to collect soft data from our TLC test platform. Based on the soft data, we simulate the performance of LDPC codes on a TLC channel using the box-plus SPA decoder [1]. Our experimental results show that, compared to hard decoding, soft LDPC decoding extends the lifetime of TLC flash memory by 4000 P/E cycles.

We then consider three detection methods and analyze their performance gain. The first method is to optimize hard decision thresholds. By applying this method to a specific page, the total number of decoding errors in one TLC block decreases from 49 to 13. The second method is to optimize multi-read threshold and an iterative algorithm is introduced to effectively find the optimal threshold. One observation emerging from the hard decision threshold optimization experiment is that the optimal hard decision thresholds varies as a function of the wordline index, resulting in an unbalanced decoding error distribution when we use the same threshold for all wordlines. The first half of wordlines in the block tend to be over-programmed and the last half of wordlines in the block tend to be under-programmed. Based on these observations, we propose the third method that separates the block into three regions and apply different thresholds to each region. By applying the third method, the total number of decoding errors in one TLC block decreases from 13 to 4.

[1] W.E. Ryan and S. Lin, *Channel Codes: Classical and Modern*, Cambridge, UK: Cambridge University Press, 2009.

NAND Flash Characterization Platform

Presenter: ***Navya Sree Prem***, PhD Student, Electrical and Computer Engineering

Collaborators: ***Industry Partner***

Advisor: ***Paul H Siegel***, Professor, Electrical and Computer Engineering

The talk will briefly go over the work towards obtaining the characterization capabilities on 3D NAND flash which would open doors to multiple avenues for further research on understanding and learning error patterns. The initial work mainly focuses on the platform for communicating with the flash on a system and processing various flash level functions. This platform eases the process of program cycling the flash, obtaining error rates, distribution and various other characteristic information.

Finding and Fixing Performance Pathologies in Persistent Memory Software Stacks

Presenter: ***Juno Kim***, PhD student, CSE

Researchers: ***Jian Xu***, PhD, Google (work done while at UCSD), ***Amirsaman Memaripour***, PhD student, CSE

Collaborators: ***Jian Xu***, PhD, Google (work done while at UCSD), ***Amirsaman Memaripour***, PhD student, CSE

Advisor: ***Steven Swanson***, Professor, CSE

Emerging fast, non-volatile memories will enable systems with large amounts of non-volatile main memory (NVMM) attached to the CPU memory bus, bringing the possibility of dramatic performance gains for IO-intensive applications. This paper analyzes the impart of state-of-the-art NVMM file systems on some of these applications and explores how those applications best leverage the performance that NVMMs offer.

Our analysis leads to several conclusions about how systems and applications should adapt to NVMMs. We propose FiLe Emulation with DAX (FLEX), a technique for moving file operations into user space and show it and other simple changes can dramatically improve application performance. We examine the scalability of NVMM file systems in light of the rising core counts and pronounced NUMA affects in modern systems, and propose changes to Linux's virtual file system (VFS) to improve scalability. We also show that adding NUMA-aware interfaces to an NVMM file system can significantly improve performance.

9500 Gilman Drive
La Jolla, California 92093
www.ucsd.edu