

Computer Science Colloquium Series since 1997

Topics, Speakers, and Abstracts during 2010-2015

- **April 3, 2015**

Compiler Optimizations for Performance and Energy Using a Combinatorial Approach

Dr. Ghassan Shobaki

Department of Computer Science

Princess Sumaya University for Technology

Abstract A compiler translates a program written in a high-level programming language into machine code. In addition to performing the translation, the compiler applies many optimizations to the generated code to improve its performance. Since nearly all contemporary programs (both applications and systems) are written in high-level languages, the compiler plays an important role in determining performance. Furthermore, a program's energy consumption is highly dependent on its performance, which means that the compiler plays an important role in minimizing energy consumption as well. Some of the important algorithmic problems involved in compiler optimizations are NP-hard problems. Examples of such problems are register allocation (mapping program variables into CPU registers) and instruction scheduling (finding the instruction order that optimizes performance). Current production compilers solve these hard problems using heuristic approaches that are not guaranteed to produce optimal solutions. My research over the past decade has focused on developing precise algorithms for solving hard compiler optimization problems using a rigorous combinatorial optimization approach, thus emphasizing the interaction between systems and theory. In this talk, I will describe the combinatorial algorithms that I have developed, with the help of my students, for solving two NP-hard compiler optimization problems: instruction scheduling for balancing instruction-level parallelism (ILP) and register pressure (the demand for CPU registers) and instruction scheduling for minimizing switching energy. Instruction scheduling with the objective of balancing ILP and register pressure is a fundamentally important problem in code generation and optimization that still does not have a satisfying solution. Due to the inherently conflicting requirements of maximizing ILP and minimizing register pressure, the heuristic techniques used by current production compilers may give poor results in many cases. The algorithm described in this talk is the first combinatorial algorithm that is efficient enough to optimally solve large instances of this problem (basic blocks with hundreds of instructions) within a few seconds per instance. The proposed algorithm uses branch-and-bound enumeration with a number of powerful pruning techniques to efficiently search the solution space. We have integrated our implementation of the proposed algorithm into the LLVM Compiler and evaluated it using SPEC CPU 2006 on real x86-64 hardware. With a time limit of 10 ms per instruction, the algorithm optimally schedules 76% of the hot basic blocks in FP2006. Another 20% of the blocks are not optimally scheduled but are improved in cost relative to LLVM's heuristic. This improves the execution time of some benchmarks by up to 22%, with a geometric-mean improvement of 2.8%. The second problem addressed in this talk is instruction scheduling to minimize the switching energy on the bus. We have formulated this problem as a Precedence-Constrained Traveling Salesman Problem

(PCTSP), which is also known as the Sequential Ordering problem (SOP), and developed an exact algorithm for solving it using branch-and-bound enumeration. The SOP has applications in other domains such as operations research. We have implemented our proposed algorithm into the LLVM Compiler and evaluated it using MiBench for the ARM target. The results show that with a time limit of 10 ms per instruction, the proposed algorithm optimally solves 99.8% of the instances. It optimally solves instances with up to 598 instructions within a few seconds. The resulting switching cost is 16% less than that produced without energy awareness and 5% less than that produced by a commonly used heuristic.

- **April 1, 2015**

Big Data at Extreme Scales
Dongfang Zhao
Department of Computer Science
Illinois Institute of Technology

Abstract In the era of Big Data, state-of-the-art yet decades old architecture of extreme-scale high performance computing systems has shown limits for many data-intensive scientific applications: compute and storage resources are separated and every I/O needs to be transferred via the network between both cliques. In this talk I will present our multi-year effort to address the I/O bottleneck of this architecture. The key idea is to introduce a distributed storage layer local to the compute nodes that is responsible for most of the I/O operations and saves extreme amount of data movement between compute and storage resources. We have designed and implemented a system prototype of such architecture, the Fusion distributed file system (FusionFS), to support metadata-intensive and write-intensive operations, both of which are critical to the I/O performance of large-scale applications. FusionFS has been deployed and evaluated on up to 16K compute nodes in an IBM Blue Gene/P supercomputer at Argonne National Laboratory, showing more than an order of magnitude performance improvement over other popular file systems such as GPFS, PVFS, and HDFS. FusionFS also employs unique designs on cooperative caching, erasure coding, transparent compression, and many more.

- **March 30, 2015**

Online, Unsupervised System Anomaly Management in Virtualized Cloud Environments
Daniel J. Dean
Department of Computer Science
North Carolina State University

Abstract Ensuring satisfactory application performance in multi-tenant cloud environments is a challenging task. When a problem missed during testing manifests in the production environment as a performance anomaly (e.g., SLO violation), there can be significant financial implications for all parties involved. In this talk I will discuss my work on predicting and diagnosing these types performance anomalies in virtualized cloud environments. First, I will present UBL, a tool for black-box performance anomaly prediction in virtualized cloud environments. UBL is designed to predict when a performance anomaly will occur by monitoring black-box metrics (e.g., CPU usage, memory usage) using an unsupervised artificial neural network called the self

organizing map. UBL is able to make these predictions while imparting negligible runtime overhead to the server. I will then discuss PerfScope, a tool which provides developers with a ranked list of suspicious functions for inspection online using a robust function-level profile. This profile is generated offline using an unsupervised data mining technique called frequent episode mining. PerfCompass and PerfScope are able to perform their analysis while imparting negligible runtime overhead to the server.

- **March 20, 2015**

Data Mining in Social Networks
Dr. Wen (Wendy) Xu
Department of Computer Science
The University of Texas at Dallas

Abstract Social networks, which consist of social actors and relationships between them, are now popular communication platforms for the public. The talk will introduce important properties about social networks, such as six degree of separation, three degree of social influence, etc. Also, it will talk about how to discover useful knowledge using data mining techniques based on the data that was collected from social networks. For instance, how to find "Kate"s (influential customers) for product promotion based on the weighted maximum cut framework, how to spot where Web rumors originate from using a concept of set resolving set, etc. Future directions about big data exploration in social networks will be discussed.

- **March 18, 2015**

Big Data Integration: Complexity and Numerosity Challenges
Dr. Wensheng Wu
Department of Computer Science
UNC Charlotte

Abstract Data integration is a long-standing challenging problem in the database community and has been extensively researched in the past several decades. In this talk, I will discuss the challenges in large-scale data integration. In particular, I will examine the "bigness" of the scale from two dimensions: the complexity of data sources and the number of data sources to be integrated. I will present our works on discovering topical and referential structures of databases that facilitate the understanding and integration of complex databases, e.g., containing hundreds or thousands of tables. I will also describe a query template-driven pay-as-go-you paradigm that enables an incremental integration of a large number of Web databases, i.e., the Deep Web, based on user query interests. Finally, I will outline the vision of a social and mobile Deep Web and possible directions to achieve the vision.

- **March 16, 2015**

Data Management for Business Processes
Yutian Sun
Department of Computer Science

Abstract A business process (BP) is a collection of activities and services assembled together to accomplish a business goal. Business process management (BPM) refers to the management and support for a collection of inter-related business processes, which has been playing an essential role in all enterprises. Business practitioners today face enormous difficulties in managing data for BPs due to the fact that the data for BP execution is scattered across databases for enterprise, auxiliary data stores managed by the BPM systems, and even file systems. Moreover, current data and business process modeling approaches leave associations of persistent data in databases and data in BPs to the implementation level with little abstraction. Implementing business logic involves data access from and to database often demands high development efforts.

In the current study, we focus on formal approaches that model the BP behavior and its associated data, capture its running status relevant to database updates, and maintain its connection with the database. A data mapping language is introduced aiming to bridge BP data and enterprise database, so that the BP designers only need to focus on business data instead of how to manipulate data by accessing the database. With the data in BPs clearly defined and mapped, we present how business data can be managed for execution that is independent from BP engines by conceptualizing and wrapping all needed data used in a BP throughout its execution. This principle enables "business process as a service", where BP execution engines or BP providers should be free of managing any client data while clients only need to manage their data instead of understanding the BP executions.

- **March 13, 2015**

Regular-Expression Search over Encrypted Data in the Cloud

Dr. Moshen Amini Salehi

Center of Advanced Computing Studies (CACS)

University of Louisiana Lafayette

Abstract Lack of trust has become one of the main concerns of users who tend to utilize one or multiple Cloud providers. Trustworthy Cloud-based computing and data storage require secure and efficient solutions which allow clients to remotely store and process their data in the Cloud. User-side encryption is an established method to secure the user data on the Cloud. However, using encryption, we lose processing capabilities, such as searching, over the Cloud data. I will present RESeED, a system that provides capability to process regular-expression based search over encrypted data residing in multiple Clouds. RESeED improves upon current state-of-the-art techniques in the search over encrypted data, by providing a highly scalable and accurate solution with low storage and performance overhead. The proposed solution is user-transparent, and cloud-agnostic. To support searching encrypted data over large-scale (and potentially big data scale) data-sets, we also have developed a scalable version of RESeED that works based on the MapReduce model.

- **March 2, 2015**

Cyber Forensics

Rahner James

Technical director
Mantech International, Cyber Security

Abstract The field of cyber forensics will be exploding over the next five years. Between state-sponsored malware and independent hackers the problems that will need to be solved will be huge with only select few that will have the skills to defend against these threats. Currently, over 250,000 new pieces of malware are submitted to a single aggregator every day. This number will grow. Many more are being created that are currently not being detected – purposefully designed to elude current signature based scanning software. The presentation will cover the following aspects: current state of the industry, tools and skills required for an analyst, current approaches to forensics, behavioral forensic analysis and new statistical heuristics concerning the flavor of data and code.

- **November 24, 2014**

Flowchart Software for the Instruction of Beginner Programmers
Devin Cook
Department of Computer Science
California State University Sacramento

Abstract Whether it was with BASIC, Pascal, C, or Java, every student had to struggle with learning their first programming language. In the past, most programming languages had a gentle learning curve. But, with the advancement of library-dependent and object-oriented languages, the rudimentary "Hello, world!" program has increased in size and complexity. As a result, the learning curve for students, rather than diminishing over time, has increased.

This project, "Flowgorithm", makes use of visual programming language concepts to ease students into a common high-level language. It makes use of industry-standard flowcharts to teach concepts such as variables, control flow and code modularization. Once students understand these concepts, and are ready to advance beyond flowcharts, Flowgorithm aids the transition using interactive code generation. Flowgorithm can generate code in C++, C#, Java, JavaScript, Lua, Python, Ruby, VBA, and Visual Basic .NET. Flowgorithm was developed at Sacramento State University.

- **October 20, 2014**

An Analysis of the Buffer Overflow Problem
Dr. Matt Bishop
Department of Computer Science
UC Davis

Abstract Significant work on vulnerabilities focuses on buffer overflows, in which data exceeding the size of a variable or an array is loaded into memory. This causes variables and state information located adjacent to the intended destination in memory to change. Unless the process checks for these additional changes, the process acts incorrectly, often placing the system in a non-secure state. We present a taxonomy based upon preconditions that must hold for an exploitable buffer overflow to exist. We validate the approach by examining several

software and hardware countermeasures. We then discuss generalizing this approach to other vulnerability types such as input errors.

This is joint work with Prof. Sophie Engle (University of San Francisco), Dr. Sean Whalen (University of California at San Francisco), and Damien Howard.

- **May 2, 2014**

A New Security Building Block and Its Application in Distributed Systems

Dr. Mehrdad Nojoumian

Department of Computer Science

Southern Illinois University Carbondale

Abstract Secret Sharing is widely used in distributed secure systems as a cryptographic primitive. In this scheme, a secret is divided into shares to be distributed among a set of players. A subset of players can then cooperate to recover the secret. In this talk, we initially propose a mathematical model of trust in social networks. Subsequently, the notion of social secret sharing is introduced in which shares are allocated based on a player's reputation and the way she interacts with other parties. Afterwards, we show that this new primitive can be used in secure systems distributed over the cloud to guarantee the service level agreement. Finally, we illustrate some ongoing research agendas.

- **April 28, 2014**

Facing a Sea of Sensed Data in Enterprise Cyber Security Defense

Jun Dai

School of Information Technology

Penn State University

Abstract Cyber security is nowadays facing “a sea of sensed data”, especially in an enterprise environment. All the information technologies in enterprise security deployment, such as firewalls and the various intrusion alert systems, contribute their data in disparate format to the sea. Detecting and preventing intrusions in cyber space is like “catching big fishes in the sea”. Like fishermen, we need a well-knit “fishing net” to capture cyber-attacks. As researchers, we first propose an interconnected enterprise-level and multiple-layer model called **SKRM** (Situation Knowledge Reference Model). This model seamlessly integrates different abstraction layers of cyber situation knowledge in an enterprise network, and serves as the fishing net. It breaks the “isolation” in-between heterogeneous data sources, and enables a “big picture” to deliver macroscopic perspective and holistic understanding.

Using SKRM as a reference model helps identify and solve problems that were previously ignored or not well addressed. Two such problems are recognized: the zero-day attack path problem and the network service dependency discovery problem. We then present solutions to them as two capable systems. One system is named **Patrol** to correctly and efficiently identify zero-day attack paths. The other one is called **Snake** to accurately and quickly discover the network service dependencies. Both systems' design benefits from the cross-layer paradigm of

SKRM. At the same time, both systems contribute to the whole construction of SKRM-enabled big picture.

In addition to Patrol and Snake, we are further inspired to build a network-wide hypervisor-based information flow monitoring and policy enforcement platform, named **X-ray**. X-ray incorporates but goes beyond Patrol and Snake: it can well fit the environment of a large-scale network and distributed system, like the cloud. Moreover, X-ray has high potentials to enable some promising capabilities. One of them, also a short-term research goal, is to extract the business work flows based on mining network service dependencies (outcome of Snake) in the real-time network-wide system object dependency graph (outcome of Patrol). The resulted business work flow graph will form another layer of SKRM and complement its construction. Our long-term goal is to fully and automatically implement SKRM, and study how it facilitates the analysts to diagnose the enterprise environment.

- **April 23, 2014**

Smart Graphics for the Ordinary User
Dr. Amit Shesh
Department of Computer Science
Illinois State University

Abstract In this talk, I will summarize why I am interested in computer graphics, and within it, the particular research problems that target various ordinary computer users. I will talk about specific research problems that I have been involved with in the last decade. They will lead to a few ideas that I am currently working on, and specific connections to my role as a prospective faculty member at CSU Sacramento. The talk will focus on visual research results (videos, demonstrations, etc.) with technical details as needed, and will target expert and budding computer scientists, knowledgeable in graphics or otherwise.

- **April 16, 2014**

Running Real-time Embedded Tasks Using GPU Computing
Pinar Muyan-Ozcelik
Department of Computer Science
UC Davis

Abstract The advent of GPU computing allows us to leverage the computation power of GPUs in different domains beyond games and graphics. Since GPUs provide significant speedups in runtime and superior cost/power-per-performance, they are a good fit for embedded domains and mobile platforms running real-time data-parallel tasks. However, meeting real-time constraints with limited hardware resources on the parallel architecture of GPUs is a nontrivial research challenge. An even more challenging research question is developing approaches for allowing GPUs to concurrently run multiple real-time embedded tasks. In this presentation, I aim to address these challenges by first presenting my study exemplifying the suitability of GPUs for performing embedded real-time tasks. Then, I go one step forward and present my research

demonstrating that GPUs can be used for running multiple disparate real-time embedded tasks at the same time.

In my first study, I perform speed-limit-sign recognition in real-time using the limited hardware resources of a low-end GPU. To meet the real-time performance requirements, I use computer vision techniques well-suited for the parallel GPU architecture, allowing me to achieve a very fast runtime, and build my pipeline from parameterized modules, allowing the best use of the limited resources by fine-tuning the parameters based on a trade-off between runtime and success rate. Since mobile gaming and graphics applications also involve several computer vision tasks that need to make the best use of limited mobile power, techniques and modules developed in the speed-limit-sign recognition study are also applicable to mobile platforms.

Like automotive computing field, mobile platforms involving game and graphic tasks that increasingly interact with the physical world require concurrently running several real-time data-parallel tasks. Hence, in my second study, considering the specific characteristics of the GPU workloads, I survey a wide spectrum of scheduling strategies for multitasking among real-time embedded tasks. Based on this investigation, I design several schedulers that use alternative approaches and determine, for a given workload, which scheduling approach is more effective and why. I also highlight the current GPU architecture shortcomings with regard to running multiple real-time tasks, and recommend new features that, when added to upcoming architectures, would allow better schedulers to be designed.

I conclude my presentation by discussing future research and teaching directions that involve developing cross-platform support for efficiently accessing GPU compute in mobile platforms running graphics and gaming applications.

- **April 14, 2014**

Big Data Privacy and Security Issues in the Clouds: The State of the Practice,
Challenges, Opportunities, and Future Directions

Dr. Arash Nourian
MIT

Abstract Big Data is arguably one of the pillars in driving advances in the next-generation of information technology. Big Data promises a better world by changing everything it touches ranging from business to private lives. It is also considered as a tool in addressing the real implications of today's society such as predicting disease patterns, creating better services, and public goods.

While the promise of Big Data is real, security and privacy issues are also magnified with the influx in the amount of data being stored in the Clouds. Also, the increasing ability for companies and governments to gain an inside look at citizens' personal lives through Big Data can fuel privacy concerns with the lack of proper security and privacy enhancing algorithms in place.

This talk will present the current state of the practice in addressing the security and privacy risks posed by Big Data. In particular, I will explain traditional security mechanisms, show their limitations, and discuss opportunities for future research.

- **March 5, 2014**

Internal workings of a solid state drive: a firmware perspective

Scott Fryman

Principle Engineer and firmware architect

PMC-Sierra

Abstract This talk will discuss what goes on inside a solid state drive from a firmware perspective. We will cover the basics of how NAND flash works and how it is used. Then will continue into how an IO is processed with respect to NAND flash. Finally, we will cover several of the main firmware algorithms required to make a solid state drive work, including garbage collection, wear leveling, and data scrubbing.

- **October 23, 2013**

DreamSpark

Doug Holland

Architect at Microsoft

Abstract Microsoft DreamSpark is designed to inspire and empower the next generation of computer science students as well as provide students and faculty with access to the latest tools and technologies.

Microsoft Windows reaches more than one billion consumers worldwide and with that reach comes an awesome opportunity for developers. Combined the Windows and Windows Phone Stores provide developers with an excellent way to build apps for the PC, tablets, and phones. In this talk Doug Holland, an architect and author at Microsoft, will show how to DreamSpark is empowering students worldwide. With DreamSpark students can access the latest versions of Windows and Visual Studio. You'll also discover how easy it is to share code between Windows and Windows Phone versions of your apps. You'll also learn of the success of others such as Kevin Ashely (a peer of Doug's at Microsoft) who made over \$100,000 with his apps on Windows and now gives away thousands of dollars to charities with his Windows Phone app!

- **October 21, 2013**

IBM Academic Initiatives

Nancy Knight

Manager, State Government Relations at IBM

- **May 1, 2013**

The Skinny on Software-Defined Networks

Chuck Black

Co-Founder and Principal Software Architect

Tallac Networks

Abstract This session will examine software-defined networking in detail. Perhaps you have heard the acronym and some of the hype, this session goes deeper. We will attempt to answer

the questions: What it is, where it started, how it works, why we care, and who is playing in the game.

- **April 24, 2013**

Datacenters and Server Virtualization
LaMills Garrett
Chief Technology Officer
HP Storage-Americas
Hewlett Packard Company

Abstract Server virtualization has become essential to most IT departments as a result of the reduced costs and consolidation of resources that's enabled by merging multiple applications into virtual servers hosted on a single or multiple physical servers. Along with the benefits of server virtualization, IT departments continue to seek other areas to reduce cost while simultaneously maximizing the benefits of server virtualization. In this colloquium, we evaluate key trends and technologies being leveraged by IT departments to help companies balance the need to drive cost lower while simultaneously enable the business to exceed competitors and find new markets. Some of the areas explored will be:

1. product designs for servers, storage, and networking,
2. datacenter co-location and datacenter design,
3. power and cooling considerations,
4. management and monitoring of systems, and
5. the impact of cloud computing.

In addition to exploring these areas, guidance will also be offered on careers, locations, and trends that will continue the evolution of computing technologies and datacenters.

- **April 10, 2013**

Journey to the Cloud
LaMills Garrett
Chief Technology Officer
HP Storage-Americas
Hewlett Packard Company

Abstract As the rage continues around Cloud and “x” as a Service, companies struggle with the journey they will plot in adopting Cloud into their business model at a time when many of them vaguely understand what “Cloud” is. Similar to the shift from mainframe to open systems brought about by the prevalence of client/server environments decades ago, Cloud stands to make a momentous shift in information technology that will greatly change how businesses leverage people, processes, and technology. This colloquium will first delve into the underlying technologies and factors driving adoption of Cloud and transition to the key questions and considerations that many businesses will face along their journey. In addition to addressing issues of technology and business, guidance will also be offered on evolving the mindsets of those that work with, purchase, and depend on technology in businesses.

- **March 6, 2013**

Data Storage Explosion and the Emergence of SSD Technologies
LaMills Garrett
Chief Technology Officer
HP Storage-Americas
Hewlett Packard Company

Abstract Based on median projections of several analysts, more data will be created in the next 12-18 months than has been created during the entire human history. The creation and storing of data is an obvious offshoot of computing technology. Since the emergence of Mainframes in the 1960s, computing technology has consistently become a pragmatic replacement for manual means of organizing, storing, and making decisions on data. Some of the key factors in inducing this data explosion are access, social connections, technology innovations, business competition, and regulations. This colloquium focuses on how technology innovations (specifically compute and storage) have contributed to the data storage explosion and how the emergence of SSD technology will further encourage the data explosion that is being witnessed today. The impact of the other factors will also be explored with an emphasis on access and social connections and how all of the factors will continue to drive larger amounts of data and lead to technologies much more integrated in our daily lives.

- **November 14, 2012**

New Enterprise Computing in the Mobile and Cloud Era
Wenjing Chu
Distinguished Technologist
Dell's Office of the CTO

Abstract "Phones Mean Freedom", that's the caption used by Pres. Bill Clinton in a recent TIME Magazine cover story listing (mobile computing) technology as the No. 1 of the "5 Ideas That Are Changing the World (for the better)". And Cloud Computing has been compared to mainframe computers and Personal Computers that will define the new paradigm of computing technology. This talk looks at what these two intertwined tectonic forces are reshaping Enterprise Computing Infrastructure and introduces some exciting technology developments that we are involved with: BYOD, end-to-end mobility, 5G Wi-Fi, SDN, Cloud and Compositional Computing.

- **October 31, 2012**

An overview of research and development (R&D) projects from the Computing Department
Dr. Alvaro Rodrigues Pereira Junior
Universidade Federal de Ouro Preto, Brazil

Abstract In this talk professor Álvaro R. Pereira Jr. will present an overview of the main projects running in the Department of Computing (DECOM) of Universidade Federal de Ouro Preto (UFOP), Minas Gerais, Brazil. The objective of the talk is to make an introduction of the

UFOP's professors, their research areas of interest, and their projects, so that professors from the California State University will be able to identify opportunities for collaboration. For instance, UFOP has professors working in several different areas of the computer science, like: computer networks, optimization, computational modeling and simulation, mobile computing, cloud computing, high-performance computing, digital libraries, image processing, signal processing, software engineering, information retrieval and web, recommender systems, machine learning, data mining, intelligent transportation systems.

- **March 26, 2012**

A Fuzzy Active Rule System for XML Databases

Dr. Ying Jin

Department of Computer Science

California State University Sacramento

Abstract Traditional database systems usually handle precise and well-defined data. In reality, there is data that is uncertain or ambiguous, where information cannot be defined precisely. Incorporating fuzzy logic into traditional crisp database systems allows database users to specify imprecise real-world semantics. Active database rules, also named Event-Condition-Active rules, have mainly been used in relational database systems for integrity control and are promising features for XML databases. This research designed a framework of fuzzy active rule system on top of existing XML databases to allow specifications of integrity constraints and business semantics. The system allows storage of fuzzy data and incorporates fuzzy expressions into standard XQuery query language. The system also includes a language framework and execution model for active rules, where fuzzy expressions can be used to express rules. We also support rule execution that is triggered by composite events and temporal events.

- **November 2, 2011**

Learning through Overcoming Inconsistency: A Step toward Perpetual Learners

Dr. Du Zhang

Department of Computer Science

California State University Sacramento

Abstract One of the long-term research goals in machine learning is how to build never-ending learners. The state-of-the-practice in the field of machine learning thus far is still dominated by the one-time learner paradigm: some learning algorithm is utilized on data sets to produce certain model or target function, and then the learner is put away and the model or function is put to work. Such a learn-once-apply-next approach is not adequate in dealing with many real world problems and is in sharp contrast with the human's life-long learning process. On the other hand, learning can often be brought on through some stimulus such as inconsistent circumstances. In this talk, we describe a framework for never-ending or perpetual learning agents where episodes of learning are triggered by inconsistency and learning is embodied in the process of overcoming inconsistent circumstances. The proposed framework facilitates recognition of inconsistency in data, information, knowledge, or meta-knowledge, and allows an agent's beliefs to be revised or augmented to explain, resolve, or accommodate inconsistency in a perpetual manner. To

substantiate the viability of the framework, we offer several specific algorithms for the inconsistency-induced learning, or i^2 Learning, which we believe is an important step toward building perpetual learning agents.

- **October 5, 2011**

Stereoscopy with Neural Network Hardware
Dr. Scott Gordon
Department of Computer Science
California State University Sacramento

Abstract Radial basis neural networks are becoming increasingly popular for computer vision. A hardware implementation produced by Cognimem, Inc (right here in Folsom, CA) is being used in a variety of real-world applications, from detecting glass manufacturing imperfections, to bin-packing on Nordic fishing boats. This talk will show how these neural network boards are being used for real-time learning and stereoscopic tracking.

- **September 28, 2011**

Information Assurance Overview
Joseph Mettle
Technical Director
National Agency Remote and Deployed Operations (RDO) Group

Abstract Mr. Mettle will have give a presentation and a video. The title of the Presentation is "Information Assurance Overview". The presentation addresses the role at the National Security Agency to protect communication networks from Cyber attacks, employment opportunities, Internships and COOPs. The video is of a recent cyber attack and defend competition between U.S. military academy universities and the National Security Agency. The universities attempted to harden their networks using their Blue Teams against a Cyber attack from the Red Teams.

- **April 27, 2011**

Feature Mining and Pattern Recognition in Steganalysis
Mengyu Qiao
Department of Computer Science and Engineering
New Mexico Tech University

Abstract Steganography, the ancient art for secretive communications, has revived on the Internet by way of hiding secret data, in completely imperceptible manners, into a digital file. Thus, steganography has created a serious threat to cyber security due to the covert channel it provides that can be readily exploited for various illegal purposes. Therefore, there is a critical need to develop reliable methods for steganalysis (analysis of multimedia for steganography) to serve applications in national security, law enforcement, cybercrime fighting, digital forensics, and network security, etc. This talk will discuss current research status of steganalysis, introduce

comprehensive performance measures, and present novel methods to defeat information-hiding in digital media.

- **April 20, 2011**

Network Security in the Core Internet
Dr. Martin Nicholes

Abstract Inter-domain routing is critical for connecting networks together to create the Internet. Local networks are under the control of a single administration, and so they employ consistent security and routing methodologies. These local networks or domains are called Autonomous Systems (ASes). Routers on the edges of these domains communicate using an inter-domain routing protocol in order to share routes. Since border routers on either side of inter-domain links are normally under different administrations, many routing problems occur, such as routing security issues, and routing control issues.

Dr. Nicholes will cover research on collaborative inter-domain routing techniques that exploit a router's position in the Internet. For example, one router could collaborate with another router to create a routing solution meeting some type of special need, such as avoiding another AS. He will present a protocol, called Inter-Domain Collaborative Routing (IDCR), to allow routers in different viewpoints to discover one another and collaborate to create routes. He presents measurements of the effectiveness of three algorithms as they are used to select ASes to deploy Inter-Domain Collaborative Routing (IDCR). In addition, he discusses results of a study of the various properties of competition among servers offering a routing service, such as the fraction of route requests handled by a service provider and the fraction of total revenue obtained. He seeks to uncover the effects on the overall network, as service providers compete. He will end with a brief discussion on Information Assurance and Virtual Machines.

- **April 6, 2011**

Impact of Battery Life on "Idle Software"
Manuj Sabharwal
Software Solution Group, Intel Corporation

Abstract Software plays an important role in battery life. Operating System, firmware, drivers and all small components are optimized to give better performance and battery life. Within the past year, shipments of notebook computers have outpaced shipments of desktops. As the notebook PC (and smaller form factor devices including tablets and smart phones) become pervasive compute platforms, battery life is becoming increasingly important, particularly stand-by or idle time. In addition, as the hardware power states become more sensitive, software must be well behaved at idle so that it doesn't needlessly wake components and limit battery life. This presentation presents two case studies that show how software "idle" behavior can have a negative impact on battery life on Windows* Operating System. It also describe a methodology to find the root cause of the behaviors and the steps necessary to mitigate the behaviors.

- **Feb. 23, 2011**

Green Computing
Abhishek Agrawal
Software Services Group, Intel Corporation

Abstract Power is increasingly becoming a central issue in designing systems, from embedded systems to data centers. The amount of energy consumed by data centers has doubled in recent years and minimizing power consumption is one of the primary technical challenges that today's IT organizations face. At the same time, the utility of mobile devices including laptops, netbooks, smartphones, tablets etc. is severely constrained by their battery life. This presentation will cover various aspects of power aware computing and energy usage in systems ranging from tablets/smartphones to data centers. The presentation will emphasize on the role of software in achieving overall platform energy efficiency.

- **Nov. 29, 2010**

Active Rule Base Verification and Analysis
Dr. Xiaoou Li
Department of Computer Science
National Polytechnic Institute (CINVESTAV-IPN), Mexico

Abstract Active rule-based systems (active systems) respond automatically to events that are taking place inside or outside the system. The most important element of active systems is the rule base which represents knowledge about a particular area. Developing a rule base is commonly perceived as a difficult task for a variety of reasons; first, it is not obvious which parts of an application should be supported using active mechanisms; second, what performance penalty is likely to result from the use of rules; and third, the verification and analysis tools which detect errors and determine if the rule base properly fulfills its goal may be minimal.

Many works can be found on production rule base verification. However, active rule base verification is much more complicated due to complex rule representation and execution. To the best of our knowledge, there is very few work reporting active rule base verification and analysis. In this talk I will introduce our recent research results on active rule base verification and analysis. Our main results are as follows:

1. Basic definitions on structural errors in active rule base, such as redundancy, inconsistency, incompleteness and circularity.
2. Original definitions on potential errors which are not currently errors, but may be errors when certain conditions and database state are fulfilled.
3. Formal methods to detect errors and potential errors using a Petri nets approach.
4. Analysis on confluence and termination properties using Petri nets.
5. Software development of tool ECAPNVerifier which permits rule base verification and analysis be done automatically.

The originality of our work is due to the formal error definitions and corresponding conceptions; without these conceptual definitions one cannot do a complete verification of active rule base. The great advantage of our Petri nets based verification and analysis method is that we can draw conclusion about errors and fulfillment of confluence and termination properties with no need to test all the rules pairs in the rule base.

- **November 3, 2010**

Scholarship for Service Program
Dr. Behnam Arad
Department of Computer Science
California State University Sacramento

Abstract The Department of Computer Science at California State University, Sacramento has recently received the NSF Cyber Security Scholarship for Service (SFS) award. The main objective of the SFS is to increase the number of information assurance and computer security professionals to meet the ever-growing demands in these fields. The award provides means to CSUS to grant scholarships to undergraduate and graduate students who agree to pursue a degree in the information assurance fields and to serve in the Federal Government in a covered position (i.e., a position in the information assurance field) for a period of up to two years upon completing academic degree requirements. This presentation provides an overview of the SFS program at CSUS including topics such as SFS history, objectives and expected outcomes, management and administration activities, impacts, and intellectual merits.

- **October 6, 2010.**

Smart Grid Cyber Security Issues
Dr. Isaac Ghansah
Department of Computer Science
California State University Sacramento

Abstract The smart grid will use automated meters, two-way digital communications technology, and advanced sensors to save energy, improve electricity efficiency and reliability. Use of these systems exposes the electrical grid to potential cyber security and privacy risks. For instance, there have been media reports of fears that a hacker could gain control of thousands, even millions, of meters and shut them off simultaneously; or a hacker might be able to dramatically increase or decrease the demand for power, disrupting the load balance on the local power grid and causing a blackout.

This seminar will discuss smart grid security and privacy issues with the view to answering the following key questions among others:

- What are some potential cyber security vulnerabilities of the smart grid?
- What cyber security best practices that can be used to handle these vulnerabilities?
- What are some potential areas of research and development (R&D) with respect to R&D cyber security of the smart grid?
- Given that smart grid is a critical infrastructure, how are government, industry, and academia working to deal with smart grid cyber security issues?
- What are some unique customer privacy issues introduced by the smart grid?