



Queen's
Graduate
Computing
Society
Conference

**The Third Annual
Queen's Graduate Computing Society Conference
(QGCSC 2012)**

**Queen's University
Kingston, ON**

May 9 – 10, 2012

Conference Summary

Schedule of Events

Wednesday May 9th, 2012

2:30 p.m. – 3:00 p.m.	Opening Remarks	Walter Light Hall Room 205
3:00 p.m. – 4:00 p.m.	Opening Keynote <i>Dr. Steve Easterbrook</i>	Walter Light Hall Room 205
4:00 p.m. – 6:00 p.m.	Programming Competition	Walter Light Hall Room 310
	Informal Networking	Grad Club
6:00 p.m. – 8:00 p.m.	Conference Social	Grad Club

Thursday May 10th, 2012

9:00 a.m. – 9:30 a.m.	Arrival, Set Up, Registration & Refreshments	Biosciences Atrium
9:30 a.m. – 10:30 a.m.	Morning Keynote <i>Dr. Esin Kiris</i>	Biosciences Atrium
10:30 a.m. – 11:30 a.m.	Track I – Student Talks BioSci 1102	Track II – Student Talks BioSci 1103
11:30 a.m. – 12:30 p.m.	Poster Session	Biosciences Atrium
12:30 p.m. – 1:30 p.m.	Lunch Break	Biosciences Atrium
1:30 p.m. – 2:30 p.m.	Panel Discussion <i>Social Media in Computer Science: Is it Here to Stay, or a One Hit Wonder?</i>	Biosciences Atrium
2:30 p.m. – 3:30 p.m.	Track I – Student Talks BioSci 1102	Track II – Student Talks BioSci 1103
3:30 p.m. – 4:30 p.m.	Final Keynote <i>Dr. Selim Akl</i>	Biosciences Atrium
4:30 p.m. – 5:00 p.m.	Closing Remarks & Awards	Biosciences Atrium

Keynote Speakers

Dr. Steve Easterbrook

*Department of Computer
Science, University of
Toronto*



Biography:

Steve Easterbrook is a professor of computer science at the University of Toronto. He received his Ph.D. (1991) in Computing from Imperial College in London (UK), and was on the faculty at the School of Cognitive and Computing Science, University of Sussex from 1990 to 1995. In 1995, he moved to the US to lead the research team at NASA's Independent Verification and Validation (IV&V) Facility in West Virginia, where he worked on software verification for the Space Shuttle and International Space Station. He moved to the University of Toronto in 1999. His research interests range from modeling and analysis of complex systems to the socio-cognitive aspects of team interaction, including communication, coordination, and shared understanding in large software teams. His most recent research focuses on how we model the earth systems that underpin our current understanding of climate change, and how scientific knowledge of climate systems is shared with non-scientists.

Abstract:

"Computing the Climate – the Evolution of Climate Models"

Detailed projections of future climate change are created using sophisticated computational models that simulate the dynamics of the atmosphere, oceans and carbon cycle. These models have evolved over the last 60 years along with scientists' understanding of the climate system. In this talk, I'll trace the history of climate modeling, from the early ENIAC weather simulations created by von Neumann and Charney, through to today's Earth System Models, in which atmosphere, ocean, vegetation and ice sheet models are coupled together to study interactions and feedbacks across the climate system. I'll also provide an overview of how the models are tested, and where the remaining uncertainties are. I'll end with a look at how the models are used in the assessments of climate change performed by the IPCC, and a glimpse at some new results that will be used in the IPCC's sixth assessment report, due to be published in 2014.

Dr. Esin Kiris

*Sr. Director,
Product Design &
User Experience
Architect,
CA Technologies*



Biography:

Esin Kiris is Principal Product Design Architect of CA Technologies. Since 1990, she has been doing Computer-Human Interaction (CHI) research and design within the software industry, Air Force, FAA, NASA, and academics. Since 1995, she has been focusing on the enterprise software design.

When Esin joined the CA User Experience (UX) team in 2004 it was still in its early stages, and her extensive experience in the UX field made her a natural co-leader of this new team. She educated and evangelized User Centered Design (UCD) principles among CA product management and development teams rapidly.

She currently leads the innovative product design initiatives of CA Technologies at Center of Excellence for Wireless and Information Technology (CEWIT), State University of New York in Stony Brook, NY USA.

Esin started her career in 1990 focusing on Human Factors and CHI research and design working within the software development industry, Air Force, FAA and NASA. Since her first job with Bell Labs in 1995, she has been working on enterprise application design and development areas. Some of her publications have been cited in several research articles and included in both the ACM Digital Library and Google Books.

Abstract:


"Customer-Focused and Innovative Software Development"

Since the start of the economic crisis all organizations have experienced huge challenges in maintaining customer relationships. We have seen the emergence of a new post-crisis customer with a clear focus on brand integrity, value of money, and quality of service. All organizations now need to deliver a more efficient, customer-focused and innovative offering than ever before to reconnect with their customers. To achieve this goal, many software development organizations have had to change their development process. Now, ongoing customer engagement, such as in the Agile development process, is key during software development. As a result, Human Computer Interaction has come to the forefront as never before in the history of software development field, creating a high demand for user experience and design (UXD) professionals.

In this talk we will discuss the two factors requiring increased attention from software developers: Customer and Innovation. Customer engagement and creating innovative software products are the goals of software development organizations in this era. To achieve these goals, the organizations face some challenges such as processes, desired employee skills and experience, and quality and speed to market. As a UXD, I will be sharing my experience, skills, projects, processes and others, to illustrate what a future employee of a software development organization might expect to encounter. I encourage students to come prepared with questions to facilitate an interactive session during my talk.

Dr. Selim Akl

*Director
School of Computing,
Queen's University*



Biography:

Selim G. Akl is a professor at Queen's University in the School of Computing, where he leads the Parallel and Unconventional Computation Group. His research interests are primarily in the area of algorithm design and analysis, in particular for problems in parallel computing and unconventional computation.

He is currently Director of the Queen's School of Computing.

Abstract:

"Nature Computes"

Look around you: Nature seems to be saying *Sum ergo computo*. In this talk I will offer some observations in support of Nature's claim of *I am, therefore I compute*.

Social Media in Computer Science: Is it Here to Stay, or a One Hit Wonder?

As part of the conference, we will be holding a Panel Discussion on the effects of social media within computer science, mainly questioning its staying power. We will hear from several different perspectives on the matter, and hopefully stir up discussion with the audience.

Panelists

Dr. James R Cordy

*School of Computing,
Queen's University*



Biography:

Jim Cordy is Professor and past Director of the School of Computing at Queen's University, Kingston, Ontario, Canada. In 1985 Dr. Cordy co-founded Holt Software Associates (HSA), a Toronto-based company specializing in educational software systems, and from 1995 to 2000 he was vice president and chief research scientist at Legasys Corporation, a software technology company specializing in legacy software system analysis and renovation.

Dr. Cordy is a founding member of the Software Technology Laboratory at Queen's University. From 1991 to 1997 he led the Software Design Technology project, a multi-university research project in software architecture research funded by the Information Technology Research Centre (ITRC, now CITO, an Ontario government Centre of Excellence). As project leader Dr. Cordy was winner of the 1994 ITRC Bank of Montreal Innovation Excellence Award and the 1995 ITRC Chair's Award for Entrepreneurship in Technology Innovation.

Wendy Powley

*School of Computing,
Queen's University*



Biography:

I have worked as the Research Associate for the Database Systems Lab, headed by Dr. Pat Martin since 1992. Our primary focus is Autonomic Computing which involves developing systems that are self-managing, self-configuring, self-optimizing and self-healing. Our system focus has been on database management systems and we are now moving into the area of Cloud computing. We have numerous conference and journal publications.

Dr. Steve Easterbrook

*Department of Computer
Science, University of
Toronto*



Biography:

Steve Easterbrook is a professor of computer science at the University of Toronto. He received his Ph.D. (1991) in Computing from Imperial College in London (UK), and was on the faculty at the School of Cognitive and Computing Science, University of Sussex from 1990 to 1995. In 1995, he moved to the US to lead the research team at NASA's Independent Verification and Validation (IV&V) Facility in West Virginia, where he worked on software verification for the Space Shuttle and International Space Station. He moved to the University of Toronto in 1999. His research interests range from modeling and analysis of complex systems to the socio-cognitive aspects of team interaction, including communication, coordination, and shared understanding in large software teams. His most recent research focuses on how we model the earth systems that underpin our current understanding of climate change, and how scientific knowledge of climate systems is shared with non-scientists.

Dr. Scott Grant

*School of Computing,
Queen's University*

**Biography:**

Scott Grant is a post-doctoral researcher with the School of Computing at Queen's University, Kingston, Ontario, Canada. His research interests include software engineering, program comprehension, topic modelling, and data mining. Scott spent four years in industry, working as a software engineer at Electronic Arts and at Google. He is the president and founder of the Queen's Game Developers Club.

Dr. Esin Kiris

*Sr. Director,
Product Design &
User Experience Architect,
CA Technologies*

**Biography:**

Esin Kiris is Principal Product Design Architect of CA Technologies. Since 1990, she has been doing Computer-Human Interaction (CHI) research and design within the software industry, Air Force, FAA, NASA, and academics. Since 1995, he has been focusing on the enterprise software design.

When Esin joined the CA User Experience (UX) team in 2004 it was still in its early stages, and her extensive experience in the UX field made her a natural co-leader of this new team. She educated and evangelized User Centered Design (UCD) principles among CA product management and development teams rapidly.

She currently leads the innovative product design initiatives of CA Technologies at Center of Excellence for Wireless and Information Technology (CEWIT), State University of New York in Stony Brook, NY USA.

Esin started her career in 1990 focusing on Human Factors and CHI research and design working within the software development industry, Air Force, FAA and NASA. Since her first job with Bell Labs in 1995, she has been working on enterprise application design and development areas. Some of her publications have been cited in several research articles and included in both the ACM Digital Library and Google Books.

Panel Moderator**Karolina Zurowska**

*School of Computing,
Queen's University*

**Biography:**

Karolina is a PhD student in the School of Computing and a member of the Modeling and Analysis in Software Engineering (MASE) Group, under the supervision of Dr Juergen Dingel. Karolina's work focuses on analysis of UML-RT state machines through symbolic execution.

Karolina is a very involved member of the School of Computing, currently serving as the President of the Graduate Computing Society, and Conference Chair for QGCSC 2012.

Student Talks

Session 1

Track I – BioSci 1102

Dynamic election based sensing and routing in Wireless Sensor Networks
Sharief Oteafy

Intelligent Context-Aware Services (iCAS)
Abdulahman Abahsain

Localization and Location Management: A Future Vision for Internet of Things
Walid Ibrahim

Track II – BioSci 1103

Incremental Test Case Generation for UML-RT Models Using Symbolic Execution
Eric Rapos, Juergen Dingel

Mutation Testing and Event Processing Language (EPL)
Lorena Guti rrez-Madro al

Automated Performance Analysis of Ultra Large Scale Systems
Haroon Malik

Session 2

Track I – BioSci 1102

Improving Group Exercise with Exergames
Tad Stach

Managing Data-intensive Workloads in Clouds
Rizwan Mian

Pose Encoding of a C-arm Fluoroscope with Accelerometer
Thomas Wolff

Track II – BioSci 1103

Clinical Decision Support Software for Prediction of Postoperative Atrial Fibrillation Following Bypass Surgery
Geoffrey Seaborn

What We Found on Our Way to Building a Classifier: A Critical Analysis of the AHA Screening Questionnaire
Quazi Abidur Rahman, Sivajothi Kanagalingam, Aurelio Pinheiro, Theodore Abraham, and Hagit Shatkay

Student Posters

Row I

Incremental Test Case Generation for UML-RT Models Using Symbolic Execution

Eric Rapos, Juergen Dingel

Analysis of error sources in electromagnetic tracking for ultrasound-guided needle navigation

Irene Ayukawa

Using Side Channels to determine Co-Residence in the Cloud

Michael (Misiu) Godfrey

IDSaaS: Intrusion Detection System as a Service in Public Clouds

Turki Alharkan, Patrick Martin

Ontology Building from Web Data

Paul Geesaman

Tech Pack on SLAs in the Cloud

Amal Khalil

A Courier-Based Integrated RFID-WSN Architecture for IoT Applications.

Ashraf Alfagih

Row II

Intelligent Transportation Systems

Dina Al-Yaseen

Mutation Testing and Event Processing Language (EPL)

Lorena Gutierrez Madroñal

Computational Ultrasound Volume Reconstruction Using a Hole Filling Algorithm with a Variable Kernel Size

Thomas Vaughan, Andras Lasso, Gabor Fichtinger

Automated Performance Engineering: Speeding the Analysis of Load tests

Haroon Malik

Contextualized Semantic Analysis of Web Services

Scott Grant, Doug Martin

Towards a Framework for Migrating Web Applications to Web Services

Asil Almonaies

Student Talks Abstracts

Dynamic election based sensing and routing in Wireless Sensor Networks

Sharief Oteafy

In a world of interconnected devices and sensors, decentralized protocols offer high adaptability to topology changes prominent in Wireless Sensor Networks (WSN). Protocols resilient to topology changes stemming from nodes dying, being added, relocating or duty cycling, improve network performance in terms of lifetime and percent of events sensed and reported. Topology dependant protocols, such as cluster-based, face many hindrances especially in terms of scalability, dynamicity, and adapting to varying traffic rates. Accordingly, a novel approach is introduced in sensing, whereby a single node is elected to report a sensed event, in a decentralized manner, thereby avoiding redundant reports by other nodes which exhaust network resources. Election is based on the node with the highest likelihood of successfully reporting the event. This protocol is coupled with a localized multi-hop routing protocol, to route that report back to the sink, by electing the most reliable next-hop neighbor to relay the report. Simulation results demonstrate the increase in network lifetime, detection/reporting efficiency, and resilience to varying node density.

Intelligent Context-Aware Services (iCAS)

Abdulrahman Abahsain

It is widely accepted that the emergence of wireless networks over the past decades has opened the door for new generations of networks where the dream of anywhere communication will become a reality. It took the use of Computers from obscurity to ubiquity. With that, the amount of information that can be provided is enormous. In light of the mobility characteristics that are common between mobile devices and wireless networks, network protocols need to take into account mobility in their design. The ultimate goal for such protocols is to be aware of the user's environment. In this talk I will explain the importance of the Context-Awareness in the field of Wireless Network and I will give a proposed mobile service that used Context in order to make the best connection.

Localization and Location Management: A Future Vision for Internet of Things

Walid Ibrahim

One of the Internet of Things challenges is to localize and manage the position of Things while they are moving.

However, different things have different capabilities and properties i.e. limited power, memory, and communication power. The intent of this paper is to discuss the major issues with localization and location management in Internet of Things.

Incremental Test Case Generation for UML-RT Models Using Symbolic Execution

Eric Rapos, Juergen Dingel

Model driven development (MDD) is on the rise in software engineering and no more so than in the realm of real-time and embedded systems. Being able to leverage the code generation and validation techniques made available through MDD is worth exploring, and is a large area of focus in academic and industrial research. However given the iterative nature of MDD, the evolution of models causes test case generation to occur multiple times throughout a software modeling project. Currently, the existing process of regenerating test cases for a modified model of a system can be costly, inefficient, and even redundant. Given this inefficiency, we aim to reuse existing test cases to avoid any unnecessary generation. It is our goal to generate test cases incrementally based on an existing test suite and some

model evolution performed on a given model. Working specifically with UML-RT state-machines, we aim to solve this problem by incrementally generating the test cases each time a model evolution step occurs. Instead of regenerating an entire test suite we will examine the existing test cases and determine which tests need to be updated, removed or whether or not we need to add more tests. Additionally, it is our goal to achieve an improved understanding of the impact of typical state machine evolution steps on test cases, and how this impact can be mitigated by reusing previously generated test cases. By determining how evolution steps impact execution and test case generation, we will be better able to understand how models evolve over time, and how this evolution can be supported efficiently. We are also aiming to implement this in a software prototype to automate and evaluate our work.

Mutation Testing and Event Processing Language (EPL) **Lorena Gutiérrez-Madroñal**

To achieve the quality of an implemented software, mutation testing has been found to be effective and popular. Unfortunately, there is not enough effort in the literature for performing mutation testing of real-time applications that process large event streams. To obtain meaningful information, it is common to see the presence of event processing queries among real-time applications. These queries are partially similar to traditional SQL queries, but provide the facilities to express rich features (e.g., pattern expression). A mistake while designing a query intended for processing event streams may result in bad user experience, abnormal program behavior, and lost business opportunities. Moreover, queries can be generated dynamically with inputs not filtered properly. This requires a tester to test the behavior of queries in presence of malicious inputs.

The QRST and UCASE research groups are collaborating to propose mutation testing to assess queries of real-time applications. We choose Event Processing Language (EPL) as our case study and develop necessary mutation operators and killing criteria to generate high quality input event streams and malicious inputs. Our proposed operators inject faults in specific features of EPL queries (pattern expression, windows of length and time, batch processing of events). Moreover, we develop operators to test SQL injection that may alter EPL query structures, as well as inject common implementation faults in clauses, operators, and null usage. The proposed operators and mutant killing criteria can be used to test some common quality issues of real-time applications such as event notification failure, delay of event reporting, unexpected and incorrect event sequences.

One of the main drawbacks of mutation testing is the computational cost involved in the execution of the large number of mutants. To reduce the execution time several different techniques have been proposed. Some of these techniques are based on the reduction of the number of generated mutants, and so we have: mutant sampling, mutant clustering, selective mutation, high order mutation, and evolutionary mutation.

Recently, the UCASE research group has developed the MuBPEL tool to perform mutation testing of Web Services Business Process Execution Language 2.0. It can be used to evaluate the quality of a test suite by checking if it can tell apart a mutant from the original program. The UCASE has also developed the Evolutionary Mutation Technique (EMT), which generates and executes only one subgroup of all the mutants. The mutants are selected by means of genetic algorithms, without loss of effectiveness in the test. Moreover, they have previously published a study called GAmEra, a system for the automatic generation of mutants for WS-BPEL compositions. GAmEra uses the EMT and can generate, execute, evaluate and classify mutants according to quality. GAmEra works in such a way that it can support any programming language for the automatic generation of mutants. It incorporates a genetic algorithm which only selects high quality mutants, thus reducing the computational cost which would be involved if all the mutants are executed. We present an architecture to generate mutation operator for EPL automatically. The architecture has three parts: original program analysis, the mutant

generator (based on GAmara) and the mutant runner. In this way, the EMT testing technique could be used with this new language, and it could be new way of checking its effectiveness.

Automated Performance Analysis of Ultra Large Scale Systems

Haroon Malik

Performance analysts rely heavily on load testing to measure the performance of their applications under a given load. During the load test, analyst strictly monitor and record thousands of performance counters to measure the run time system properties such as CPU utilization, Disk I/O, memory consumption, network traffic etc. The most frustrating problem faced by analysts is the time spent and complexity involved in analyzing these huge counter logs and finding relevant information distributed across thousands of counters. We present our methodology to help analysts by automatically identifying important performance counters for load test and comparing them across tests to find performance gain/loss. Further, our methodology help analysts to understand the root cause of a load test failure by finding previously solved problems in test repositories. A case study on load test data of a large enterprise application shows that our methodology can effectively guide performance analysts to identify and compare top performance counters across tests in limited time thereby archiving 88% counter data reduction.

Improving Group Exercise with Exergames

Tad Stach

Sedentary activities have led to a decrease in the fitness level of Canadians over the last two decades. This trend is concerning since sedentary lifestyles have been associated with numerous health risks such as cardiovascular disease, high blood pressure, and osteoporosis. Physical activity can improve health and fitness; however, people face many psychological barriers to exercise. Exergames have been proposed as a means of encouraging physical activity by combining exercise and entertainment. Multiplayer exergames attempt to take advantage of the motivating aspects of group activity by allowing two or more people to exercise together. In most multiplayer exergames, a player's performance is limited by her physical abilities. This situation makes it difficult for people of disparate physical abilities to play exergames together. I will present a novel player balancing technique to better support engaging experiences in multiplayer exergames. More specifically, heart rate monitoring is used to scale in-game performance relative to how closely a person adheres to her target heart rate. I will discuss the results of two user studies evaluating the benefits and limitations of this technique.

Managing Data-intensive Workloads in Clouds

Rizwan Mian

The promise of "infinite" resources given by the cloud computing paradigm has led to recent interest in exploiting clouds for large-scale data-intensive computing. Given this supposedly infinite resource set, we need a management function that regulates application workload on these resources. This doctoral research focuses on two aspects of workload management, namely scheduling and provisioning. We propose a novel framework for workload execution and resource provisioning, and associated models, algorithms, and protocols.

Pose Encoding of a C-arm Fluoroscope with Accelerometer

Thomas Wolff

Purpose: Intra-operative X-ray machines are being used in numerous medical interventions. In these devices the X-ray generator and detector is typically mounted on a C-shaped arm that allows free rotation around two orthogonal axes. One common interventional procedure is prostate brachytherapy.

The goal of brachytherapy is to eliminate cancer by implanting radioactive capsules (seeds) in the tumor. The success of the procedure depends on how accurately the seeds are inserted in the intended locations. Algorithms exist to determine seed positions from multiple X-ray images, however they require accurate information about the C-arm's pose (i.e., position and orientation). In an operating room obtaining the pose is challenging, as clinical requirements limit the applicability of commonly used tracking methods such as optical tracking, electromagnetic tracking, image-based tracking with fiducials, or built-in mechanical encoders. Methods: We propose an alternative method to recover the C-arm's pose using an accelerometer. The accelerometer is mounted on the mobile X-ray generator. When the C-arm is stationary, the measured acceleration direction corresponds to the gravitational force direction, therefore it can be used as a tilt sensor. The accelerometer was calibrated as a tilt sensor with respect to the C-arm's rotation along its two axes, using a high-accuracy optical tracker as a reference. The scaling and offset error of the sensor was compensated using polynomial fitting. The system was evaluated in an experimental setup that mimicked the clinical procedure of brachytherapy seed position verification with fluoroscopy imaging. Results: The accelerometer-based pose measurement error for Primary Angle rotation was -0.12° (STD= 0.67°) and for Secondary Angle rotation it was -0.25° (STD= 0.14°). The built-in mechanical encoder-based pose measurement error for Primary Angle rotation was -0.04° (STD= 0.71°) and for Secondary Angle rotation it was -0.30° (STD= 0.28°). Conclusion: The accelerometer provided similar or more accurate results than the built-in mechanical encoder, thus accelerometer is a good candidate for pose tracking whenever mechanical encoders are not available or not accurate enough. The accelerometer recovered the C-arm's pose information with acceptable accuracy for prostate brachytherapy seed position computation.

Clinical Decision Support Software for Prediction of Postoperative Atrial Fibrillation Following Bypass Surgery **Geoffrey Seaborn**

Every year, approximately 11,000 Ontarians undergo coronary artery bypass grafting (CABG), which is an invasive surgical procedure designed to treat coronary artery disease (CAD), the leading cause of cardiovascular deaths in Canada. Following this critical procedure, survival in patients at high risk is greatly increased. However, many CABG recipients develop sustained atrial fibrillation (AF), the most frequent basis for hospitalizations due to heart rhythm disturbances, and the primary cause of stroke. Postoperative AF is a tremendous burden for patients in terms of morbidity and mortality, as well as for the healthcare system in terms of additional costs incurred. Despite investigators reporting numerous independent predictors for postoperative AF, its etiology continues to be unclear, and its prevention and management remain suboptimal. Efforts to reduce the incidence of postoperative AF would achieve significant health benefits for patients, and substantial cost savings for the Canadian healthcare system.

CABG is a common procedure performed principally to relieve symptoms of angina and to reduce risk of death from CAD. Conduits from elsewhere in the patient's vasculature are grafted to the coronary arteries to bypass atherosclerotic vessels, improving the blood supply to the circulation supplying the heart muscle. This surgery is typically performed with the heart stopped, necessitating the use of a cardiopulmonary bypass machine. Although the value of CABG has been extensively demonstrated, it carries risk for adverse complications. Sustained postoperative AF is the most common side effect of CABG, with onset generally occurring 2-4 days into recovery.

AF, an abnormal heart rhythm characterized by chaotic electrical activity in the atria and, consequently, by irregular contraction intervals in the ventricles, may result in palpitations, hypotension, pneumonia, syncope, chest pain, as well as congestive heart failure (CHF). AF is associated with a 3- to 5-fold increase in stroke risk. It is estimated that 350,000 Canadians are currently living with AF, and its prevalence is predicted to increase. Numerous approaches may be used to treat AF, such as synchronized electrical cardioversion, catheter ablation, and medications that either

slow heart rate or convert heart rhythm back to normal. Individuals with AF can be prescribed anticoagulants, such as warfarin, to decrease stroke risk.

Incidence of sustained AF following CABG is generally 20%-40%. The length of hospital stay after surgery is significantly longer for patients who exhibit AF than for those who do not, corresponding to considerable hospital charges. In addition, at four to five years, survival is worse in patients who develop postoperative AF than in those who do not. Investigators have reported independent predictors for its onset, such as advanced age, hypertension, resting heart rate of greater than 100 beats per minute, chronic heart failure, male gender, obesity, smoking, long signal-averaged P-wave (SAPW), P-wave dispersion, increased atrial ectopy, variations in autonomic tone, and low heart rate variability (HRV). Well-tolerated preventative measures have also been described.

The problem of accurately predicting the onset of sustained postoperative AF in patients undergoing CABG remains open. Investigators have reported many clinical indices currently associated with postoperative AF following CABG. Contemporary machine learning techniques are well-suited to recognizing underlying trends in 'training' data consisting of several labeled examples, and to using the results to classify new unlabeled data with remarkable sensitivity and specificity. We are currently developing an advanced clinical decision support software capable of automatically gathering and analyzing relevant clinical data from patients undergoing CABG in order to provide physicians with objective and non-invasive insights into the likelihood of sustained postoperative AF so that patient morbidity and mortality, as well as healthcare costs, can be significantly reduced by targeting appropriate preventative therapies.

What We Found on Our Way to Building a Classifier: A Critical Analysis of the AHA Screening Questionnaire

Quazi Abidur Rahman, Sivajothi Kanagalingam, Aurelio Pinheiro, Theodore Abraham, and Hagit Shatkey

Inherited cardiovascular disease is the main cause of sudden cardiac death (SCD) in young athletes. In the United States the incidence has been reported as 1:50,000 – 1:100,000 per year. Initial screening through electrocardiogram (ECG) and echocardiogram (Echo) is a standard first step for identifying morphological anomalies that can lead to potential cardiac abnormalities, and in extreme cases - to sudden death. However, ECG and Echo are not always used for large-scale screening of all young athletes because of the associated cost and perceived impracticality of these procedures.

As a preventive measure, the American Heart Association (AHA) has recommended a screening procedure, intended as a cost-effective and practical initial measure for pre-participation screening of athletes. The questionnaire is widely used throughout the United States as a screening tool, but concerns exist regarding the efficacy of this questionnaire in discriminating normal from non-normal heart condition.

As a component within an extensive research on the possible contribution of the questionnaire to large-scale analysis of adverse heart conditions, we set out to pursue what appeared to be a straightforward task: namely, training a machine-learning-based classifier to automatically categorize athletes into risk-levels based on their respective answers to the questionnaire. The study is based on information gathered from 470 young athletes who have followed the recommended AHA 12-element screening using the questionnaire, as well as through an additional full-fledged screening by Electro- and Echo-cardiogram examinations administered by a physician.

We expected to be able to effectively train a classifier from the questionnaire data due to the hypothesis driving the AHA guidelines, as discussed above, that the answers to the pre-screening questionnaire can indeed be correlated with the diagnosis obtained from more extensive and time consuming standard initial cardiovascular tests, (namely, Echo and ECG), administered by a physician.

Intending to follow the common machine learning procedures for learning a classifier from data, we also aimed to select the most informative features, that is, identify the components in the AHA-based pre-screening whose answers are most predictive of the cardiologist's adjudication.

This effort uncovered significant issues with the effectiveness of the questionnaire itself as a screening tool, which lead us to pursue an extensive information-content and probability-based analysis of each element of the AHA questionnaire, examining the ability of affirmative or negative answers to identify potential cardiac abnormalities.

The results of our analysis indicate that the AHA-recommended 12 elements do not effectively distinguish between normal and non-normal heart conditions and that the elements in the questionnaire are not informative of potential cardiac abnormalities, as identified by Electro- and Echo-cardiogram examinations.

Student Posters Abstracts

Incremental Test Case Generation for UML-RT Models Using Symbolic Execution Eric Rapos, Juergen Dingel

Model driven development (MDD) is on the rise in software engineering and no more so than in the realm of real-time and embedded systems. Being able to leverage the code generation and validation techniques made available through MDD is worth exploring, and is a large area of focus in academic and industrial research. However given the iterative nature of MDD, the evolution of models causes test case generation to occur multiple times throughout a software modeling project. Currently, the existing process of regenerating test cases for a modified model of a system can be costly, inefficient, and even redundant. Given this inefficiency, we aim to reuse existing test cases to avoid any unnecessary generation. It is our goal to generate test cases incrementally based on an existing test suite and some model evolution performed on a given model. Working specifically with UML-RT state-machines, we aim to solve this problem by incrementally generating the test cases each time a model evolution step occurs. Instead of regenerating an entire test suite we will examine the existing test cases and determine which tests need to be updated, removed or whether or not we need to add more tests. Additionally, it is our goal to achieve an improved understanding of the impact of typical state machine evolution steps on test cases, and how this impact can be mitigated by reusing previously generated test cases. By determining how evolution steps impact execution and test case generation, we will be better able to understand how models evolve over time, and how this evolution can be supported efficiently. We are also aiming to implement this in a software prototype to automate and evaluate our work.

Intelligent Transportation Systems Dina Al-Yaseen

Intelligent Transportation Systems (ITS) and their applications have attracted significant attention in research and industry. Existing ITS schemes depend on having the vehicles or roadside devices equipped with sensing and communication units. We propose CrowdITS that uses the information reported by the Crowd of drivers without the need of any special sensors. We integrate the Crowd information with other traffic information sources, analyze and disseminate the information to the users of the CrowdITS system based on their location.

Analysis of error sources in electromagnetic tracking for ultrasound-guided needle navigation

Irene Ayukawa

Introduction: Fluoroscopy and CT guided spinal needle injections are often used to treat chronic low back pain. However, these imaging methods are expensive and harmful to the patient and physician. Ultrasound imaging, combined with electromagnetic tracking of the needle and the ultrasound probe, offers a safe and inexpensive approach for guiding spinal interventions. Before tracked ultrasound navigation systems can be used clinically, it must be shown that they are sufficiently accurate. In this work, we investigated the accuracy of the electromagnetic tracking system.

Methodology: The accuracy was measured as a function of horizontal distance from the field generator. Two sizes of needle sensors were used to compare the performance of the system, and the experiments were performed multiple times to investigate the consistency of the results. We studied three sources of error present in our system: the needle calibration error; the needle-image registration error; and the needle placement error.

Results: Overall, the measured error values increased as the experiment is performed farther away from the field generator, as did the associated variation between results. We also found that the larger needle sensor was tracked with greater accuracy and consistency, compared to the smaller one.

Conclusion: Our electromagnetic tracking system was sufficiently accurate within accepted values. More research is still needed to show if this system can replace the current procedure for spinal interventions.

Mutation Testing and Event Processing Language (EPL)

Lorena Gutiérrez-Madrojal

To achieve the quality of an implemented software, mutation testing has been found to be effective and popular. Unfortunately, there is not enough effort in the literature for performing mutation testing of real-time applications that process large event streams. To obtain meaningful information, it is common to see the presence of event processing queries among real-time applications. These queries are partially similar to traditional SQL queries, but provide the facilities to express rich features (e.g., pattern expression). A mistake while designing a query intended for processing event streams may result in bad user experience, abnormal program behavior, and lost business opportunities. Moreover, queries can be generated dynamically with inputs not filtered properly. This requires a tester to test the behavior of queries in presence of malicious inputs.

The QRST and UCASE research groups are collaborating to propose mutation testing to assess queries of real-time applications. We choose Event Processing Language (EPL) as our case study and develop necessary mutation operators and killing criteria to generate high quality input event streams and malicious inputs. Our proposed operators inject faults in specific features of EPL queries (pattern expression, windows of length and time, batch processing of events). Moreover, we develop operators to test SQL injection that may alter EPL query structures, as well as inject common implementation faults in clauses, operators, and null usage. The proposed operators and mutant killing criteria can be used to test some common quality issues of real-time applications such as event notification failure, delay of event reporting, unexpected and incorrect event sequences.

One of the main drawbacks of mutation testing is the computational cost involved in the execution of the large number of mutants. To reduce the execution time several different techniques have been proposed. Some of these techniques are based on the reduction of the number of generated mutants,

and so we have: mutant sampling, mutant clustering, selective mutation, high order mutation, and evolutionary mutation.

Recently, the UCASE research group has developed the MuBPEL tool to perform mutation testing of Web Services Business Process Execution Language 2.0. It can be used to evaluate the quality of a test suite by checking if it can tell apart a mutant from the original program. The UCASE has also developed the Evolutionary Mutation Technique (EMT), which generates and executes only one subgroup of all the mutants. The mutants are selected by means of genetic algorithms, without loss of effectiveness in the test. Moreover, they have previously published a study called GAmEra, a system for the automatic generation of mutants for WS-BPEL compositions. GAmEra uses the EMT and can generate, execute, evaluate and classify mutants according to quality. GAmEra works in such a way that it can support any programming language for the automatic generation of mutants. It incorporates a genetic algorithm which only selects high quality mutants, thus reducing the computational cost which would be involved if all the mutants are executed. We present an architecture to generate mutation operator for EPL automatically. The architecture has three parts: original program analysis, the mutant generator (based on GAmEra) and the mutant runner. In this way, the EMT testing technique could be used with this new language, and it could be new way of checking its effectiveness.

Using Side Channels to determine Co-Residence in the Cloud

Michael (Misiu) Godfrey

Cloud computing is gaining popularity by allowing customers to outsource their IT needs. It holds efficiency over the traditional method by allowing disparate clients to share the same computing resources, while using software technologies to isolate them from one another. While many of the security challenges of using the Cloud are extensions of those found in other fields, the use of the same resources by mutually distrusting parties adds a truly unique attack surface to the paradigm; specifically, the threat that users can use qualities of the shared physical hardware to bypass the software isolation systems in place. In this poster we replicate a portion of work previously done by Ristenpart et al. to demonstrate how hardware properties can be exploited to bypass Virtual Machine (VM) isolation in the Cloud, highlighting potentially exploitable vulnerabilities. Using previous work as a guideline, we have developed a program able to communicate reliably across VM isolation through interference patterns in the use of a shared resource (the hard drive disk). We go on to demonstrate that this program functions on Amazon's EC2, a canonical implementation of a modern Cloud service.

Computational Ultrasound Volume Reconstruction Using a Hole Filling Algorithm with a Variable Kernel Size

Thomas Vaughan, Andras Lasso, Gabor Fichtinger

Purpose: Computational Ultrasound Volume Reconstruction is a process where the pixel intensities from one or more tracked ultrasound slices are inserted into voxels of an ultrasound volume. Due to inadequate sampling, however, not all voxels within the volume will be assigned pixels intensities from the original images. Such holes impact the usability and accuracy of the volume reconstruction. Holes can be filled by interpolating across voxel intensities in a local surrounding kernel region, but current algorithms will fail when the kernel region contains insufficient information. We investigated the effect of enlarging the local kernel region when this occurs.

Methods: Several series of tracked ultrasound images were collected for volume reconstruction. A hole filling algorithm that features a variably-sized kernel region was implemented, and tested in volume reconstructions of the collected data. The resulting volumes were evaluated qualitatively by visual assessment and quantitatively using the calculated Mean Absolute Error (MAE).

Results: The new hole filling algorithm filled holes that would not be filled using a static smaller kernel region. Images that implemented hole filling with a variable kernel size were easier to interpret than images that did not use hole filling. The error in hole voxels was reduced when hole filling with a variable kernel size was applied.

Conclusion: The proposed hole filling algorithm was implemented. It is capable of filling more holes than conventional hole filling algorithms, thereby making ultrasound volumes readily more usable.

IDSaaS: Intrusion Detection System as a Service in Public Clouds **Turki Alharkan, Patrick Martin**

In a public cloud computing environment, consumers cannot always just depend on the cloud provider's security infrastructure. They may need to monitor and protect their virtual existence by implementing their own intrusion detection capabilities along with other security technologies within the cloud fabric. Intrusion Detection as a Service (IDSaaS) targets security of the infrastructure level of a public cloud (IaaS) by providing intrusion detection technology that is highly elastic, portable and fully controlled by the cloud consumer. A prototype of IDSaaS is described.

Automated Performance Engineering: Speeding the Analysis of Load tests

Haroon Malik

Performance analysts rely heavily on load testing to measure the performance of their applications under a given load. During the load test, analyst strictly monitor and record thousands of performance counters to measure the run time system properties such as CPU utilization, Disk I/O, memory consumption, network traffic etc. The most frustrating problem faced by analysts is the time spent and complexity involved in analyzing these huge counter logs and finding relevant information distributed across thousands of counters. We present our methodology to help analysts by automatically identifying important performance counters for load test and comparing them across tests to find performance gain/loss. Further, our methodology help analysts to understand the root cause of a load test failure by finding previously solved problems in test repositories. A case study on load test data of a large enterprise application shows that our methodology can effectively guide performance analysts to identify and compare top performance counters across tests in limited time thereby archiving 88% counter data reduction.

A Courier-Based Integrated RFID-WSN Architecture for IoT Applications.

Ashraf Alfagih

We introduce a hierarchical architecture for the Internet of Things that utilizes ubiquitous wireless nodes in the environment as couriers between the integrated RFID-WSN components and the base-stations.

Ontology Building from Web Data

Paul Geesaman

The semantic web is an endeavor to represent knowledge in a way that is both human and machine-readable. The semantic web is made of ontologies, which represent knowledge in a field, called the ontology's domain. The building blocks of an ontology is a concept, an object or idea, and relationships between concepts.

The creation of a new ontology by hand is time-consuming. A better solution is to create a semi-automated, and unsupervised method of extracting relations between concepts. Attempts have been

made to create semi-automated engines to create new ontologies; but semi-automated is computationally heavy and can lead to missing important concepts and their relations. The process of extracting ontologies in a semi-automated manner is called 'open ontology learning'. The goal of my research is to combine OntoCMaps, an open ontology learning engine, with Cerno, a light-weight lexical analyzer, to create a more accurate and lightweight open ontology learning system.

First, an ontology for a specific domain will be built using the combination of OntoCMaps and Cerno to filter relevant information from corpora. To test the accuracy, a subject will be chosen which has a very codified set of expressions. A human will read and create a semantic web and the human's ontology will be compared to the model using measures such as F-measure. A second experiment is to verify whether the combination of the models is consistent with each other. The corpora will be split in two parts, and then the system will run each part independently to verify that the ontologies are accurate and consistent.

Contextualized Semantic Analysis of Web Services **Scott Grant, Doug Martin**

The poor locality of operation descriptions expressed in the Web Service Description Language (WSDL) makes them difficult to analyze and compare in web service discovery tasks. This problem has led us to develop a new method for service operation comparison involving contextualizing operation descriptions by inlining related type information from other sections of the service description. In this paper, we show that this contextualization of web service descriptions can enable topic models (statistical techniques for identifying relationships) to produce semantically meaningful results that can be used to reverse engineer service-oriented web systems and automatically identify related web service operations. Specifically, we model contextualized WSDL service operations using Latent Dirichlet Allocation, and show how this approach can be used to more accurately find similar web service operations.

Tech Pack on SLAs in the Cloud **Amal Khalil**

Self-manageable Cloud infrastructures are required in order to comply with users' requirements defined by Service Level Agreements (SLAs). Service providers must be capable of monitoring their infrastructure resource metrics to enforce the agreed service terms. This study is designed to offer an overview on the researches done or still ongoing toward building the infrastructure of SLA-based Cloud environment that facilitates SLA management and enforcement.









Towards a Framework for Migrating Web Applications to Web Services **Asil Almonaies**

Migrating traditional legacy web applications to web services is an important step in the modernization of web-based business systems to more complex inter-business services and interactions. While the problem of migrating various kinds of legacy software systems to a service oriented architecture (SOA) environment has been well studied in the literature, approaches to migrate dynamic web applications to web services are lacking. We will outline the requirements for a semi-automated approach to migrate dynamic legacy web applications to web services-based SOA applications while preserving the original web application's business processes. A manual demonstration of our approach is presented using two examples from Moodle, a popular open source course management system.

Programming Competition Teams

The Meerkats (Equis Lab)	Md Ameer Hamza
	Cheryl Savery
	Max Graham
	Tad Stach
Cordy's Angels (Software Technology Lab)	Doug Martin
	Paul Geesaman
	Matthew Stephan
	Andrew Stevenson
The Dream Team (Graduate Computing Society)	Eric Rapos
	Melissa Trezise
	Karolina Zurowska

Conference Organizers

Conference Chair	Karolina Zurowska	
Logistics	Eric Rapos	
Sponsorship	Gehan Selim	
Webmaster	Anton Barua	
Organizing Committee	Melissa Trezise	
	Sherin Abdel Hamid	
	Umme Hunny	
	Layan Nahlawi	

Thank You to our Sponsors!

