AGENDA

ORF-RE Research Meeting
Delivering Ultra-large-scale Services
June 19/20, 2013

Wednesday June 19, 2013
6:30 PM – BBQ & Social
Location: Grad Club, Middlesex College, Western University

Thursday June 20, 2013
Location: Middlesex College 105b, Western University

9:00 am – 9:30 am **Welcome, Introductions & Overview** Pat Martin

---

9:30 am – 10:30 am **Utilizing Mobile Ubiquitous Resources (MUR)**
Hossam Hassanein, Queen's University

A growing abundance of smart objects are largely underutilized. Everyday new smartphones enter the market with more sophisticated resources packed in multi-homed devices. At the same time, newer vehicles are being equipped with more sophisticated on board units (OBUs), encompassing an array of sensors. In both industries, the installed resources share a single attribute: they serve the user.

In this talk we will explore the rising interest in public sensing systems that attempt to utilize these abundant resources to leverage sensing and communication infrastructures. We explore the potential of utilizing vehicles as resource providers, and the power of crowd-sourced sensing via smartphones, to boost the potential of Internet of Things (IoT) paradigms. We highlight recent developments at the Telecommunications Research Lab (TRL) at Queen's University that address optimal resource utilization and public sensing paradigms geared towards synergetic solutions.

10:30 am – 11:00 am **Break**

11:00 am – 12:00 pm **Software Defect Management in Large-Scale Systems**
Mehdi Amoui, University of Waterloo (Tahvildari)

Defects in software are inevitable and impose a huge economic impact on the product. Having an efficient defect management system can have a great impact on the efficiency of enormous software ecosystems. Part of the defect management process is to understand, prioritize, and plan to resolve defects considering organizational objectives. This is known as defect triaging. Triaging is particularly important in large-scale software systems as the defect pool size can grow large and many (conflicting) objectives need to be met. This talk elaborates on the defect triaging process at BlackBerry to gain a better understanding of the shortcomings and challenges of the current practices. Throughout this talk we will: i) present our recent advances in defect investigation and the design and development of a search-based duplicate (and similar) defect detection system for BlackBerry, and ii) propose a number of practical research directions that can improve defect prioritization in large-scale software organizations.

Improving Bug Location using Correlations in Crash Reports
David (Shaohua) Wang, Queen's University (Zou)

To accelerate bug localization process, we conducted an empirical study to investigate crash correlations and their effect on improving bug localization. Using our method and findings, the effort of examining and fixing bugs can be reduced by only examining the top files ranked by our method using crash correlations.

12:00 pm – 2:00 pm **Lunch**
Grad Club, Middlesex College

1:00 pm – 2:00 pm **Posters**
Middlesex College, 105b

2:00 pm – 3:00 pm

Components with Embedded Specification -Based Security Monitors
Umair Khan, Queen's University (Zulkernine)

Security specifications have to be correctly implemented to avoid vulnerabilities in the developed software. However, some security vulnerabilities may still be present when the software is deployed. Security vulnerabilities in a software component can compound the damage as the component will be reused multiple times in different environments. Given that, a software component's behavior must be continuously observed against its security specifications after deployment to identify incorrect implementations.

To overcome the aforementioned issues, we propose that a software component should have an embedded specification-based security monitor. We first identify the activities necessary to develop a software component with an embedded specification-based security monitor. We also elaborate upon the design (constituent parts) and the working (interactions between the constituent parts and the component) of the embedded monitor. Finally, we propose security monitoring techniques to observe different constraints and behaviors imposed by security specifications.

We evaluate our proposed approach by embedding security monitors in component-based software (CBS). These CBS are compared with their counterparts that have external monitors with respect to design complexity, size, and performance. The viability of the proposed security monitoring techniques is assessed by calculating the size, memory usage, and monitoring overhead.

Test Engineering: Speeding the Analysis of Load Tests
Haroon Malik, Queen's University (Hassan)

Load testing is one of the means for evaluating the performance of Large Scale Systems (LSS). At the end of a load test, performance analysts must analyze thousands of performance counters from hundreds of machines under test. These performance counters are measures of run-time system properties such as CPU utilization, Disk I/O, memory consumption, and network traffic. Analysts observe counters to find out if the system is meeting its Service Level Agreements (SLAs). In this paper, we present and evaluate one supervised and three unsupervised approaches to help performance analysts to 1) more effectively compare load tests in order to detect performance deviations which may lead to SLA violations, and 2) to provide them with a smaller and manageable set of important performance counters to assist in root-cause analysis of the detected deviations. Our case study is based on load test data obtained from both a large scale industrial system and an open source benchmark application. The case study shows, that our wrapper-based supervised approach, which uses a search-based technique to find the best subset of performance counters and a logistic regression model for deviation prediction, can provide up to 89% reduction in the set of performance counters while detecting performance deviations with few false positives (i.e., 95% average precision). The study also shows that the supervised approach is more stable and effective than the unsupervised approaches but it has more overhead due to its semi-automated training phase.

A Concept Analysis Approach for Guiding Users in Service Discovery

Bipin Upadhyaya, Queen's University (Zou)

Web services are widely used as basic constructs to build complex distributed applications with fast speed and low cost. However, existing service discovery techniques provide users with poor results which require substantial human intervention to filter the services to locate the desired ones. In particular, users often have no prior knowledge of the functional description of the available services on the Web. The queries formulated by the users may not match well with the service descriptions of existing services. As a consequence, a user's query can result in a large number of returned services. In this paper, we propose an approach that derives the semantic concepts conveyed in the service descriptions and clusters the services based on the concepts. As a result, each concept is associated with a set of relevant services. To understand the semantic meanings of a user's query, we identify concepts behind the query and recommend related concepts associated with services. Our approach also guides users to formulate their queries. We conducted a case study and found that the average precision and recall of our approach for service discovery are respectively, 83% and 100%. We also performed a user study which shows that for 85% of time, a user reformulates their queries using the suggestion provided by our approach to improve the precision of the retrieved services.

3:00 pm – 3:15 pm	**Break**	
3:15 pm – 4:00 pm	**Next Steps & Wrap Up** Cascon Workshop Plans	Pat Martin