IBM

CASCON 2015

25th anniversary!

With a CASCON Account? Login Now!

CASCON 2015 Proceedings is Now Available for Download!

CASCON 2015 Program Guide is Now Available for Download!

CASCON 2015 Survey!

Most Influential Paper

Tracking Time-Varying Parameters in Software Systems with Extended Kalman Filters
Tao Zheng, Carleton University, Canada
Jinmei Yang, Carleton University, Canada
Murray Woodside, Carleton University, Canada
Marin Litoiu, IBM Canada Lab.
Gabriel Iszlai, IBM Canada Lab.

Best Paper

Recent News

(NEW!) CASCON 2015 Workshop Materials are now available for Download!

(NEW!) CASCON 2015 Survey is now Available

(NEW!) CASCON 2016, Nov. 8th-10th. Mark you calendar. Program co-chairs for CASCON 2016 are Hausi A. Müller, University of Victoria and Vio Onut, IBM Canada CAS

(NEW!) Recorded keynotes, FoSP sessions and workshop presentations will be available for download after the conference. Stay tuned.

Reminder: CASCON 2015 Opening Keynote Session starts at 9AM on Monday, Nov 2nd, 2015

CASCON 2015 Program Guide is available for download Now

CASCON 2015 Proceedings is available for download Now

Need help on CASCON website? Contact cwillis@ca.ibm.com or casresearch@ca.ibm.com

Follow us through these social outlets:

Facebook

Google+

LinkedIn

Twitter

Related Links

The Four Languages of Innovation: A Model of Industrial Research and Development (1MB)

Research Partner

The Effect of a Collaborative Game on Group Work
Maaz Nasir, Rockstar Games, Canada
Kelly Lyons, University of Toronto, Canada
Rock Leung, SAP Canada
Anthea Bailie, Markham Public Library, Canada
Fred Whitmarsh, Markham Public Library, Canada

Best Student Paper

Cardinality Estimation Using Neural Networks
Henry Liu, University of Waterloo, Canada
Mingbin Xu, York University, Canada
Ziting Yu, University of Waterloo, Canada
Vincent Corvinelli, IBM Canada Lab.
Calisto Zuzarte, IBM Canada Lab.

Faculty Fellow of the Year

Kostas Kontogiannis, National Technical University of Athens

Student of the Year

Michael Athanasopoulos, National Technical University of Athens

Collaborator of the Year

Piotr Mierzejewski, Software Development Manager, DB2 Build & Infrastructure, IBM Analytics

Innovator of the Year

Ontario Centres of Excellence (OCE)

Consortium for Software Engineering Research (CSER)

National Research Council Canada (NRC)

Interested in Becoming a Faculty Fellow?

Click Here to Read about the Requirements

.

Cynthia Zhang, User Experience Designer - IoT, PLE, IBM Analytics

Project of the Year

Prioritizing and Visualizing Security Findings
Dr. Lin Tan, Assistant Professor, University of Waterloo
Jinqiu Yang, Ph.D., Candidate, University of Waterloo
John Peyton, Architect, AppScan Source, IBM Security
Stephen Teilhet, SAST Security researcher, AppScan Source, IBM Security
Kristofer Duer, Lead Security Analytics Researcher, AppScan Source, IBM Security
Dr. Iosif Viorel Onut, Principal R&D Strategist, IBM Center for Advanced Studies

Best Technology Showcase

Monitoring business goals with Argus
Alejandro Mate, University of Alicante
Juan Trujillo, University of Alicante
Elvis Koci, University of Trento
Konstantinos Zoumpatianos, University of Trento
John Mylopoulos, University of Trento

People's Choice Technology Showcase

Context Awareness in Video Analytics
Nina Taherimakhsousi, University of Victoria
Hausi A. Muller, University of Victoria

CASCON 2015 Conference Theme: Redefining the New Enterprise

Transformational technologies in cloud, mobile, cognitive computing, internet of things, social computing, analytics and big data and others have become vortexes that change the enterprise landscape. For enterprises to stay relevant and competitive, facing the challenge of understanding and leveraging these new technologies is a must. Enterprises who can leverage new technologies to lead with new enterprise paradigms will set themselves apart from their competitors. They can deliver distinguishing appeals as they engage their customers. They can take the business to a newer level of success that others can't. Commerce, healthcare, finance, smarter cities are examples of industrial domains going through such technological evolution. This conference provides a forum to ideate, discuss, evaluate the depth of each transformational technology, as well as their confluence impact in redefining new enterprises.

About CASCON

CASCON was established in 1990. CASCON is a premiere industrial academic conference in computer science and software engineering, with about a thousand three hundred attendees. It continues to be an innovative forum for researchers, innovators, technologies and developers from academia, industry, and government to share knowledge, form collaborations, present original work and experience first-hand the latest technology trends.

Conference Content and format include:

- Keynote presentations
- Technical paper track
- Emerging technology track for position papers
- Workshops in different technology areas (such as Technology Areas)
- Best Paper and Best Student Paper Awards
- Frontiers of Software Practice plenaries
- Technology showcase of research outcome

CASCON proceedings are available in the ACM Digital Library and DBLP indexed.

CASCON 2015 Program Co-Chairs

The co-chairs for CASCON 2015 are :

Hanan Lutfiyya, Western University, Canada

Marin Litoiu, York University, Canada

Technical Papers

The technical papers program will feature both experience reports and original research papers in the general area of software technology and applications, but particularly addressing this year's theme.
At the conference, awards are given out for Best Paper and Best Student Paper. Criteria for the paper includes originality, clarity and potential impact. Two paper awards - for Best Paper and Best Student Paper - will be presented to recognize the best technical contributions of the event in terms of originality, clarity and potential impact. To be eligible for the "Best Student Paper" award, a paper must have been primarily authored by a student and the work that the paper reports must have been undertaken by the student authors. Only the student author(s) of the Best Student Paper award receives a prize. The recipients must have been students at the time the work was undertaken.

Call for Technical Papers, Important Dates :

Abstracts for Technical Papers Due June 26, 2015

Final Technical Papers Due July 3, 2015

Acceptance Notifications to Authors - July 31, 2015

Camera-ready Technical Papers Due August 21, 2015

Emerging Technologies Track (ETT)

CASCON Emerging Technology Track invites position papers to share visions, ideas, and insights on these advancements, highlight technology gaps, research agendas, and integrated solutions, and initiate discussion, feedback, and further collaboration.

Call for Position Papers, Important Dates :

Abstracts for Position Papers Due August 14th, 2015

Final Position Papers Due August 21st, 2015

Acceptance Notifications to Authors - September 18th, 2015

Camera-ready position papers Due October 2nd, 2015

Hotel

Hilton Suites Toronto,Markham Conference Centre & Spa

8500 Warden Avenue, Markham,Canada
L6G 1A5

905-470-8500

IBM

Conference Schedule

Below is the conference schedule for CASCON 2015.

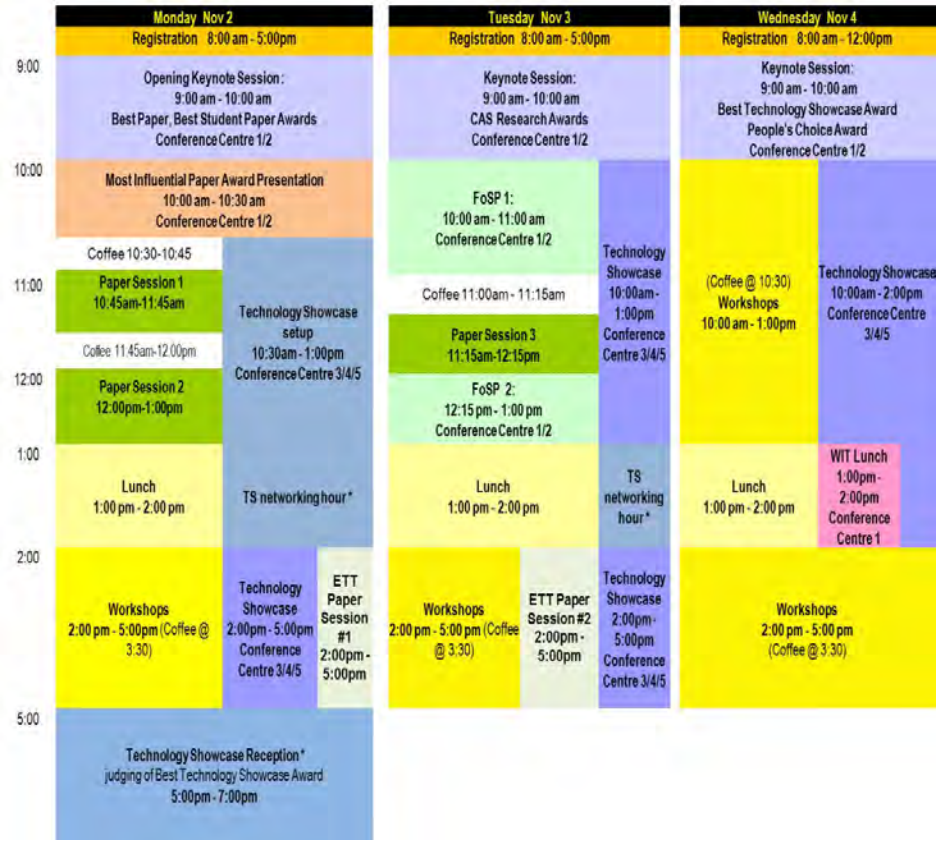

Recent News

(NEW!) CASCON 2015 Workshop Materials are now available for Download!

(NEW!) CASCON 2015 Survey is now Available

(NEW!) CASCON 2016, Nov. 8th-10th. Mark you calendar. Program co-chairs for CASCON 2016 are Hausi A. Müller, University of Victoria and Vio Onut, IBM Canada CAS

(NEW!) Recorded keynotes, FoSP sessions and workshop presentations will be available for download after the conference. Stay tuned.

Reminder: CASCON 2015 Opening Keynote Session starts at 9AM on Monday, Nov 2nd, 2015

CASCON 2015 Program Guide is available for download Now

CASCON 2015 Proceedings is available for download Now

Need help on CASCON website? Contact cwillis@ca.ibm.com or casresearch@ca.ibm.com

Follow us through these social outlets:

Facebook

Google+

LinkedIn

Twitter

Related Links

The Four Languages of Innovation: A Model of Industrial Research and Development (1MB)

Research Partner

Ontario Centres of Excellence (OCE)

Consortium for Software Engineering Research (CSER)

National Research Council Canada (NRC)

Interested in Becoming a Faculty Fellow?

Click Here to Read about the Requirements

.

IBM

Keynote Speakers

CASCON 2015 will feature three keynote sessions as well as two FoSP sessions. All sessions take place in the Conference Center 1 and 2.

To view speakers' bio please click on their respective images.

Time: 9:00 AM - 10:00 AM EST, Monday, November 2nd
Adam Kocoloski
DE & CTO, Cloud Data Services, IBM Analytics

Data Services in the Cloud: Past, Present and Future

In some respects the modern cloud computing era began with the launch of a data service, namely the S3 object store. The subsequent decade has witnessed a rapid expansion in the depth and breadth of available data services catering to the entire spectrum of workloads: structured and unstructured, transactional and analytic, interactive and archival. All this variety notwithstanding, we can identify elements of value that are common across the spectrum. Developers are increasingly embracing cloud data services as a way to speed innovation, improve IT economics at scale, and avail themselves of world-class operational expertise.

Of course, the journey of data into the cloud has enountered a bit of weather along the way. The rapid consolidation onto highly-scalable, multi-tenant systems has magnified the impact of any individual outage, and in our quest for performance and scale we have sometimes sacrificed valuable constructs borne of decades of research and experience. There are no easy answers to some of these issues, but that's what makes working in this field so much fun. We can evaluate new innovations in large-scale experiments and incorporate them into production systems on rapid timescales. With this reality in mind, we'll conclude the talk by reviewing some major research areas of interest for the Cloud Data Services group at IBM.

Recent News

(NEW!) CASCON 2015 Workshop Materials are now available for Download!

(NEW!) CASCON 2015 Survey is now Available

(NEW!) CASCON 2016, Nov. 8th-10th. Mark you calendar. Program co-chairs for CASCON 2016 are Hausi A. Müller, University of Victoria and Vio Onut, IBM Canada CAS

(NEW!) Recorded keynotes, FoSP sessions and workshop presentations will be available for download after the conference. Stay tuned.

Reminder: CASCON 2015 Opening Keynote Session starts at 9AM on Monday, Nov 2nd, 2015

CASCON 2015 Program Guide is available for download Now

CASCON 2015 Proceedings is available for download Now

Need help on CASCON website? Contact cwillis@ca.ibm.com or casresearch@ca.ibm.com

Follow us through these social outlets:

Facebook

Google+

LinkedIn

Twitter

Related Links

The Four Languages of Innovation: A Model of Industrial Research and Development (1MB)

Research Partner

Time: 9:00 AM - 10:00 AM EST, Tuesday, November 3rd
Jim Caldwell
Director, IBM Internet of Things, Continuous Engineering Solutions Development

IBM Internet of Things Point of View and Strategy

The Internet of Things is predicted to have an economic impact of more than $11 Trillion per year by 2025. It has become a focus of discussion by technologists, the business press and the general public. Clearly something is happening but what? And what should businesses and institutions do about it?

This presentation will survey the topic from IBM's perspective. We will discuss what the Internet of Things is. We will also discuss IBM's point of view and strategy, some examples of offerings and client engagements. Finally, we will conclude with some key questions and research challenges.

Time: 9:00 AM - 10:00 AM EST, Wednesday, November 4th
Dr. Eliot Siegel
Professor and Vice Chairman Research Information Systems University of Maryland School of Medicine Department of Diagnostic Radiology and Nuclear Medicine and Chief Imaging VA Maryland Healthcare System

Big Data and a Challenge to IBM to Help Reinvent How We Can Use it in Medical and Imaging Applications in Healthcare

"BIG DATA" is currently one of the hottest topics in medicine from a research and clinical perspective. However, it's impossible to get any consensus on its definition. IBM and others have defined it in terms of the four V's, volume, velocity, variety, and, particularly relevant to healthcare, veracity

Whichever definition that we may use when thinking of "Big Data", medicine and in specific, diagnostic imaging clearly generates vast amounts of it. The volume and complexity of medical information in healthcare has doubled every five years with 80% or more of that data unstructured.

One of the major challenges with medical imaging is the difficulty of discovery of imaging information in the electronic medical record and from clinical trial data. Our imaging reports are, almost without exception, unstructured and our medical images are rarely tagged in such a way as to be discoverable or useful to data mining efforts. This must change if medical imaging is to play a substantial role in this era of big data, medical guidelines, decision support and personalized medicine.

The goals of this presentation include defining "Big Data" and issues such as

Ontario Centres of Excellence (OCE)
Consortium for Software Engineering Research (CSER)
National Research Council Canada (NRC)

Interested in Becoming a Faculty Fellow?

Click Here to Read about the Requirements

.

archival of clinical images and other data and metadata, exploring the current role and applications of big data in diagnostic imaging and medicine with an emphasis on clinical applications and speculating about the potential and future applications of Big Data to help in visualization of images and data, and diagnosis and treatment. Real-time clinical applications based on real time mining of clinical trial data for decision support and screening will be presented. An emphasis will be placed on outlining the challenges we have to practical use of Big Data in medicine.

FoSP Speakers

Time: 10:00 AM - 11:00 AM EST, Tuesday, November 3rd
Dr. Frank Dehne
Chancellor's Professor of Computer Science
Director, Institute for Data Science
Carleton University, Ottawa, Canada

Analytics Performance: The Role of Parallel Analytics Computation

Analytics computation and big data often require a high preformance computing platform such a large scale cloud infrastructure or dedicated compute cluster. Parallel (multi-threaded) compution is required to leverage such large scale compute infrastructure. In some cases, without high performance computing, analytics computation is impossible. Here, we present a particular example from analytics computation in biomedical computing: the design of new molecular markers, inhibitory agents, or drugs based on inhibitory compounds that inhibit disease related cellular functions while leaving non-target functions unaffected (to minimize side-effects).

Time: 12:15 PM - 1:00 PM EST, Tuesday, November 3rd
Tim Francis
IBM Distinguished Engineer, IBM Commerce

IBM Commerce: Leveraging Data Science to Enable Customer and Partner Engagement

A famous scene in the movie "Minority Report" shows Tom Cruise walking into a Gap store in 2054; as he is biometrically recognized, the store's marketing automation systems suggest items he might like to buy. Today, a mere 12 years after the release of that movie, the world is fast converging on that vision, particularly in our digital lives. A first step in successful omni-channel customer engagement lies in curating and harmonizing a wide variety of information, such as exposure / response to marketing campaigns, online and offline purchases, price responsiveness, social media interaction, and product availability; with this foundation an enterprise can cultivate invaluable customer insights that improve both profitability and customer loyalty.

Women in Technologies (WIT) Speaker

Time: 1:00 PM - 2:00 PM EST, Wednesday, November 4th
Paul Zikopoulos
Vice President, Client Success - IBM Analytics Big Data, Competitive, and Technical Sales. IBM Analytics, Solutions

Time: 1:00 PM - 2:00 PM EST, Wednesday, November 4th
Joanna Ng
Innovator, Chief Innovation Office for Cloud & IoT, STSM and Master Inventor

WiT - Memoirs from a Nobody

Women in technology has always been a controversial topic often rising up in discussion. What progress have we made? What progress is there still to be made? Crucial issues like this should involve the community as a whole, however it's noticeably more women being involved in such conversations. It's time to break taboos and boundaries and talk about things openly. Why does the technology field have to be gender related? Instead of talking about women in technology or men in technology, why don't we contemplate about what makes a person successful?

This year, CASCON is pleased to welcome on stage Paul Zikopoulos, Vice President, Client Success, as key speaker. The speech will subsequently be followed by an interactive exercise - discussion among the audience.

IBM

Technical Papers

The technical papers program will feature both experience reports and original research papers in the general area of software technology and applications, but particularly addressing this year's theme.

At the conference, awards are given out for Best Paper and Best Student Paper. Criteria for the paper includes originality, clarity and potential impact.

Two paper awards - for Best Paper and Best Student Paper - will be presented to recognize the best technical contributions of the event in terms of originality, clarity and potential impact. To be eligible for the "Best Student Paper" award, a paper must have been primarily authored by a student and the work that the paper reports must have been undertaken by the student authors. Only the student author(s) of the Best Student Paper award receives a prize. The recipients must have been students at the time the work was undertaken.

CASCON 2015 Accepted Papers

Congratulations to the following authors on their technical papers being accepted by CASCON 2015. They will present their papers during the conference on Nov. 2nd and 3rd.

CLOUD COMPUTING and INFRASTRUCTURE/ Monday, November 2nd 10:45AM - 11:45AM
(Evergreen)
Session chair: Anwar Haque

Title: Optimizing Application Downtime through Intelligent VM Placement and Migration in Cloud Data Centers
Author: Venkatesh Nandakumar, Madalin Mihailescu, Cristiana Amza, Zartab Jamil, Alan Wen Jun Lu and Harsh Singh

Abstract for Optimizing Application Downtime through Intelligent VM Placement and Migration in Cloud Data Centers
As cloud data centres grow in size and complexity, hosted applications become increasingly vulnerable to dynamically occurring infrastructure downtime periods caused by partial infrastructure failures. Downtimes within cloud data centres can be diverse, ranging from

Relevant documents

Paper Template 2015

Cascon 2015 Copyright Form (18KB)

Recent News

(NEW!) CASCON 2015 Workshop Materials are now available for Download!

(NEW!) CASCON 2015 Survey is now Available

(NEW!) CASCON 2016, Nov. 8th-10th. Mark you calendar. Program co-chairs for CASCON 2016 are Hausi A. Müller, University of Victoria and Vio Onut, IBM Canada CAS

(NEW!) Recorded keynotes, FoSP sessions and workshop presentations will be available for download after the conference. Stay tuned.

Reminder: CASCON 2015 Opening Keynote Session starts at 9AM on Monday, Nov 2nd, 2015

CASCON 2015 Program Guide is available for download Now

CASCON 2015 Proceedings is available for download Now

Need help on CASCON website? Contact cwillis@ca.ibm.com or casresearch@ca.ibm.com

Follow us through these social outlets:

Facebook

Google+

LinkedIn

Twitter

Related Links

unplanned server/rack unit failures to compulsory server power-offs when addressing arbitrary environment conditions, e.g., thermal issues. For instance, in these environments, server racks are often a unit of failure due to either faulty rack switches or rack power units. We observe that the degree of application disruption depends on i) the application's fault tolerance, reconfiguration capabilities, and redundancy of VM components affected by the respective emergency shutdowns and ii) the support for VM migration of vulnerable application components within the constrained time window of impending shutdown of a failure unit. In this context, in this paper, we develop and evaluate techniques which aim to optimize the downtime of hosted applications during emergency shutdowns due to partial failures through two orthogonal approaches: i) designing VM placement techniques that are aware of application fail-over semantics and ii) prototyping intelligent schemes for live VM migration prioritization based on prediction models for both VM migration times and expected downtimes for applications.

Title: A Study of Three MapReduce Frameworks
Author: Kiran Sundaravarathan, Aparna Balachandra Bhat and Patrick Martin

Abstract for A Study of Three MapReduce Frameworks
Apache Hadoop is an open source platform that is used for storage and parallel processing of huge amounts of data. It has been adopted by a number of organizations. The wide adoption of Hadoop by industries has prompted users to stretch its usage beyond the map-reduce paradigms for which it was designed. This along with a number of factors has led to the development of the next generation Hadoop called Apache YARN, which is a cluster management technology that can support different programming paradigms. At the same time, Facebook discovered that the original Hadoop architecture could not handle their scalability and processing needs. Hence they modified the architecture of Hadoop to create a new framework called Corona, which shares some of the features found in YARN. In this paper, we examine and contrast the three parallel programming frameworks. We also present a set of experiments to compare Hadoop, YARN and Corona with respect to their methods and ease of deployment and their performance for well-known benchmarks.

Title: Adoop: MapReduce for Ad-hoc Cloud Computing
Author: Mohammad Hamdaqa, Mohamed Sabri, Akshay Singh and Ladan Tahvildari

Abstract for Adoop: MapReduce for Ad-hoc Cloud Computing
MapReduce is a widely adopted distributed programming model for big data analytics. It facilitates processing large data sets using a dedicated cloud containing many worker nodes. For most implementations of MapReduce to work, only a few nodes may fail at any given time. This poses a challenge for organizations trying to harness the power of underutilized computing resources by provisioning cloud services on top of existing IT infrastructure. Exploiting adhoc clouds for MapReduce operations could yield significant advantages. It could reduce operational costs, improve resource utilization, and enable big data analytics. To the best of our knowledge, only few previous researchers have tried to optimize MapReduce for such volatile, non-dedicated, environments. This paper investigates how Hadoop -the most widely used open-source implementation of MapReduce- can be optimized to run efficiently in ad-hoc cloud environments, despite the challenges these environments impose. To address these challenges, we present Adoop: a history-based scheduling approach to MapReduce, where the availability history of each node affects Hadoop scheduling decisions. Adoop maintains an availability and utilization based score of all the participating nodes, and dynamically re-adapts task assignments accordingly. A proof-of-concept implementation of Adoop has been provided and made publically available. Our initial experiments show that Adoop outperforms Hadoop in a simulated ad-hoc environment.

SOFTWARE TOOLS and PROCESSES/ Monday, November 2nd 10:45AM - 11:45AM (Jasmine)
Session chair: Kelly Lyons

Title: CrashAutomata: An Approach for the Detection of Duplicate Crash Reports Based on Generalizable Automata

The Four Languages of Innovation: A Model of Industrial Research and Development (1MB)

Research Partner

Ontario Centres of Excellence (OCE)

Consortium for Software Engineering Research (CSER)

National Research Council Canada (NRC)

Interested in Becoming a Faculty Fellow?

Click Here to Read about the Requirements

.

Author: Neda Ebrahimi Koopaei and Abdelwahab Hamou-Lhadj

Abstract for CrashAutomata: An Approach for the Detection of Duplicate Crash Reports Based on Generalizable Automata

Crash reporting systems are useful tools that allow users to report system failures, subsequently contacting the appropriate support group for resolution. As a software system grows and becomes more versatile, the number of crashes increases. A large software company receives typically thousands of crashes a day, which make it difficult for software engineers to address these reports in a timely manner. Fortunately, not all reports are new; many of them are duplicates of previously reported crashes. Research has shown that early detection of duplicate reports can reduce the effort and time it takes to handle crash reports. In this paper, we propose a new approach for detecting duplicate crash reports, called CrashAutomata. CrashAutomata builds a model from historical crash reports (more precisely their stack traces) that is used to classify an incoming report. The model is based on varied-length n-grams and automata. Unlike existing techniques, CrashAutomata takes advantage of the generalization aspect of automata, making it possible to build a representative model of crash reports, reducing the number of false positives. When applied to crash reports of the Firefox system, CrashAutomata results in very high precision and recall. It also outperforms CrashGraph, a leading technique in the detection of duplicate crash reports.

Title: Towards Convenient Management of Software Clone Codes in Practice: An Integrated Approach

Author: Sharif Uddin, Chanchal K. Roy and Kevin Schneider

Abstract for Towards Convenient Management of Software Clone Codes in Practice: An Integrated Approach

Software code cloning is inevitable during software development and unmanaged cloning practice can create substantial problems for software maintenance and evolution. Current research in the area software clones includes, but is not limited to: finding ways to manage clones; gaining more control over clone generation; and, studying clone evolution and its effects on the evolution of software. In this study, we investigate tools and techniques for detecting, managing, and understanding the evolution of clones, as well as design a convenient tool to make those techniques available to a developer's software development environment. Towards the goal of promoting the practical use of code clone research and to provide better support for managing clones in software systems, we first developed SimEclipse: a clone-aware software development platform, and then, using the tool, we performed a study to investigate the usefulness of using a number clone based technologies in an integrated platform rather than using those discretely. Finally, a small scale user study is performed to evaluate SimEclipse's effectiveness, usability and information management with respect to some pre-defined clone management activities. We believe that both researchers and developers would enjoy and utilize the benefits of using SimEclipse for different aspects of code clone research as well as for managing cloned code in software systems.

Title:A Monitor-based Synchronization Approach for Java Packed Objects

Author: Bing Yang, Kenneth Kent, Eric Aubanel and Karl Taylor

Abstract for A Monitor-based Synchronization Approach for Java Packed Objects

Packed object support is an experimental enhancement for Java in the IBM J9 Virtual Machine. In this paper we present a synchronization approach for the packed object data model, which is monitor-based, but in a lockword-free manner. This approach seeks to address the distinctive synchronization issues and situations which are faced by the packed object data model and have never been addressed by traditional Java built-in synchronization mechanisms before. According to the experimental results, our approach performs much better than the Java built-in synchronization for coarse-grained locking. The time overhead is dropped by 18% for the PackedSync approach. For fine-grained locking, although the performance of our approach is overtaken by the Java built-in synchronization as array sizes and the number of threads increase, PackedSync performs reasonably similar while array sizes and thread numbers are small.

NEXT GENERATION SYSTEMS/ Monday, November 2nd 10:45AM - 11:45AM (Primrose)

Session chair: Mike Smit

Title: Stratified Sampling for Even Workload Partitioning Applied to Single Source Shortest Path Algorithm
Author: Jeeva Paudel, Levi H.S Lelis and Jose Nelson Amaral

Abstract for Stratified Sampling for Even Workload Partitioning Applied to Single Source Shortest Path Algorithm
An efficient implementation of large graph processing algorithms on distributed-memory machines requires a balanced partitioning of the graph across the machines. Recently, Paudel et al. presented an algorithm, named Workload Partitioning and Scheduling (WPS), that uses domain-specific knowledge to guide a sampling procedure in large implicitly-defined graphs. WPS's sampling procedure is used for partitioning the workload into parts of similar size which is then distributed amongst different machines. This work extends Paudel et al.'s study and presents an investigation of the parallel and distributed implementation of Meyer's Triangle-Stepping algorithm for solving the Single Source Shortest Path (SSSP) problem for directed graphs. Our implementation leverages the WPS algorithm for evenly distributing the workload involved in processing the vertices of the input graph across distributed-memory machines. In contrast with Paudel et al.'s study which focused on implicitly-defined graphs, this work demonstrates that WPS is also equally applicable on explicitly-defined graphs. Empirical evidence shows that applying WPS to Meyer's SSSP algorithm yields significant performance benefits.

Title: Benchmarking the IBM Power8 Processor
Author: Istvan Reguly and Abdoul-Kader Keita

Abstract for Benchmarking the IBM Power8 Processor
This report discusses the performance of IBM's Power8 CPUs, on a number of skeleton, financial and CFD benchmarks and applications. Implicitly, the performance of the software toolchain is also tested - the bare-bones Ubuntu, the beta XL compilers and OpenMP runtimes. First, we aim to establish some roofline numbers on bandwidth and compute throughput, then move on to benchmark explicit and implicit one-/three-factor Black-Scholes computations, and CFD applications based on the OP2 and OPS frameworks, such as Airfoil, Volna, BookLeaf, CloverLeaf, CloverLeaf 3D and Rolls-Royce Hydra. These applications all exhibit different characteristics in terms of computations, communications, memory access patterns, etc. Both absolute and relative performance metrics are computed and compared to NVIDIA GPUs and Intel Xeon CPUs.

Title: MetaFork: A Compilation Framework for Concurrency Models Targeting Hardware Accelerators and its Application to the Generation of Parametric CUDA Kernels
Author: Changbo Chen, Xiaohui Chen, Abdoul-Kader Keita, Marc Moreno Maza and Ning Xie

Abstract for MetaFork: A Compilation Framework for Concurrency Models Targeting Hardware Accelerators and its Application to the Generation of Parametric CUDA Kernels
In this paper, we present the accelerator model ofMetaFork together with the software framework that allows automatic generation of CUDA code from annotated MetaFork programs. One of the key features of this CUDA code generator is that it supports the generation of CUDA kernel code where program parameters (like number of threads per block) and machine parameters (like shared memory size) are allowed. These parameters need not to be known at code-generation-time: machine parameters and program parameters can be respectively determined and optimized when the generated code is installed on the target machine. This generation of parametric CUDA kernels requires from the MetaFork framework to deal with non-linear polynomial expressions during the dependence analysis and tiling phase of the MetaFork code. To achieve these algebraic calculations, we take advantage of quantifier elimination and its implementation in the RegularChains in Maple. Various illustrative examples are provided together with performance evaluation.

SOCIAL COMPUTING and SMARTER INTERACTIONS/ Monday, November 2nd

12:00PM - 1:00PM (Evergreen)
Session chair: Vio Onut

Title: The Effective of a Collaborative Games on Group Work
Author: Maaz Nasir, Kelly Lyons, Rock Leung, Anthea Bailie and Fred Whitmarsh

Abstract for The Effective of a Collaborative Games on Group Work
As enterprises strive to transform themselves in the face of emerging technologies and challenges, there is an increasing reliance on group work and collaboration across disciplinary and organizational boundaries. Often this requires establishing ad-hoc workgroups that come together to address a new problem over a short period of time. In these cases, it is important that group members can work together effectively in a short amount of time. We are interested in understanding how computer-based "icebreaking" video games can help group members work together in real world collaborations. To do this, we identified ice-breaking video game requirements based on the literature and ran an experiment with ad-hoc workgroups within an organization to assess the effect of playing an ice-breaking video game before one of their collaborative work tasks. We compared groups that participated in the ice-breaking video game prior to the work task with those that did not. We found that groups that played the icebreaking video game demonstrated increased collaboration in the subsequent work task.

Title: Context-Aware Mobile Apps using iBeacons: Towards Smarter Interactions
Author: Edward Sykes, Stephen Pentland and Saverio Nardi

Abstract for Context-Aware Mobile Apps using iBeacons: Towards Smarter Interactions
In this paper we describe four mobile apps for iOS devices that use Bluetooth Low Energy iBeacons to provide contextual relevance and personalized experiences for the user. The applications span a number of vertical markets including asset tracking, food transportation logistics and health care. We developed these apps in collaboration with an industry partner located in Mississauga, Ontario, Canada. In this paper we present the relevant background of work in this area, the architectural framework that we designed and developed to support these context-aware apps, the apps themselves, and report on the findings of real use test case scenarios.

MACHINE LEARNING and ANALYTICS/ Monday, November 2nd 12:00PM - 1:00PM (Jasmine)
Session chair: Sudhakar Ganti

Title: Cardinality Estimation Using Neural Network
Author: Henry Liu, Mingbin Xu, Ziting Yu, Vincent Corvinelli and Calisto Zuzarte

Abstract for Cardinality Estimation Using Neural Network
Database query optimizers benefit greatly from accurate cardinality estimation; however, this is hard to achieve on tables with correlated and/or skewed columns. We present a novel approach using neural networks to learn and approximate selectivity functions that take a bounded range on each column as input, effectively estimating selectivities for all relational operators. Experimental results with a simplified prototype show a significant improvement over state-of-the-art cardinality estimators on constructed datasets in terms of accuracy, efficiency, and amount of user input required.

Title: Automated Classification of Congestive Heart Failure Severity using Time Domain, Frequency Domain and Non-linear Heart Rate Variability Measures
Author: Mehrin Gilani and Mikael Eklund

Abstract for Automated Classification of Congestive Heart Failure Severity using Time Domain, Frequency Domain and Non-linear Heart Rate Variability Measures
Congestive Heart Failure (CHF) is one of the leading causes of death in elderly in Canada. It has a 5-year survival rate of around 50% and half a million Canadians live this disease. CHF

is a progressive disease that rapidly increases in severity. As a result, CHF patients have to pay unscheduled visits to the hospital due to critical emergencies. Lack of automated techniques for the prediction and detection of CHF not only degrades the quality of life of these patients but also causes them extreme financial stress. Automated techniques for the detection of critical events can help clinicians monitor these patients' cardiac health more efficiently. In this paper, we present an automated classifier for the detection of CHF severity. We classified New York Heart Association class I, II and III patients using time domain, frequency domain and non-linear heart rate variability (HRV) measures. We compared the performance of our multi-class classifier and the binary classifier using different sets of HRV features. Our results show that using a combined set of features instead of Standard Deviation of NN intervals (SDNN) alone, improves the classifier accuracy by almost 21%. Moreover, using HRV measures extracted from longer duration of NN intervals, improve the classification accuracy of class I in multi-class classifier by almost 3 times.

SOFTWARE ANALYSIS and MODELING/ Monday, November 2nd 12:00PM - 1:00PM
(Primrose)
Session chair: Hausi Muller

Title: An Empirical Study on Change Recommendation
Author: Manishankar Mondal, Chanchal K. Roy and Kevin Schneider

Abstract for An Empirical Study on Change Recommendation
Recommending changes to programmers by exploiting their repetition tendencies during system evolution has been investigated by a number of studies. In our research we per- form a change type (additions, deletions, and modifications) based analysis of the efficiency of change recommendation. We also investigate the programmer sensitivity of the repeated changes (i.e., the extent the same changes are repeated by the same programmers) of different change types. The existing studies did not perform such investigations. However, these investigations can be important for efficient ranking (i.e., prioritizing) and altering of recommendations. According to our investigation on thousands of commits of five diverse subject systems we observe that modifications have a very low tendency (around 1.3%) of being repeated. We should primarily focus on recommending additions, and deletions. More importantly, overall 71% of the repeated changes are programmer sensitive. We believe that a change recommendation system that prioritizes recommendations considering programmer sensitivity can help programmers reuse previous changes in a time-efficient manner.

Title: Context Extraction in Recommendation Systems in Software Engineering: A Preliminary Survey
Author: Sana Maki, Segla Kpodjedo and Ghizlane El Boussaidi

Abstract for Context Extraction in Recommendation Systems in Software Engineering: A Preliminary Survey
Recommendation Systems in Software Engineering (RSSEs) represent a new promising research area, whose goal is to help software developers in their tasks by providing them with context-dependent insights extracted from their current project or taken from best practices. A key challenge here is to retrieve the context from the programming task in order to provide useful recommendations. In this paper, we conduct a survey of RSSEs with a particular focus on different approaches used to extract the context. We propose a feature model to represent some important characteristics of such extraction and identify some open issues.

SOFTWARE MAINTENANCE and EVOLUTION/ Tuesday, November 3rd 11:15AM - 12:15PM
(Evergreen)
Session chair: Eleni Stroulia

Title: Recommending Relevant Sections from a Webpage about Programming Errors and Exceptions
Author: Mohammad Masudur Rahman and Chanchal K. Roy

Abstract for Recommending Relevant Sections from a Webpage about Programming Errors and Exceptions
Programming errors or exceptions are inherent in software development and maintenance, and given today's Internet era, software developers often look at web for finding working solutions. They make use of a search engine for retrieving relevant pages, and then look for the appropriate solutions by manually going through the pages one by one. However, both the manual checking of a page's content against a given exception (and its context) and then working an appropriate solution out are non-trivial tasks. They are even more complex and time-consuming with the bulk of irrelevant (i.e., off-topic) and noisy (e.g., advertisements) content in the web page. In this paper, we propose an IDE-based and context-aware page content recommendation technique that locates and recommends relevant sections from a given web page by exploiting the technical details, in particular, the context of an encountered exception in the IDE. An evaluation with 250 web pages related to 80 programming exceptions, comparison with the only available closely related technique, and a case study involving comparison with VSM and LSA techniques show that the proposed technique is highly promising in terms of precision, recall and F1-measure.

Title: Do Developers Respond to Code Stability Warnings?
Author: Sylvie Foss and Gail Murphy

Abstract for Do Developers Respond to Code Stability Warnings?
Ideally, developers would always release code without bugs. Given the impossibility of achieving this ideal, there has been growing interest in ways to alert a developer earlier in the development process to code that may be more bug prone. A recent study found Google developers were unsure of how to act on file-level bug prediction information provided during code reviews as developers did not know how to take action on agged files. In this paper, we investigate whether presenting direct metrics about code stability in a developer's editor may be more amenable to developer attention and action. We show that prior study results about code stability performed on telecommunications software apply to our industry partner's fast-evolving and large JavaScript system. Building on this result, we introduce the ChangeMarkup plugin we built to display code stability in terms of code age and commit sizes via per-line markers in the code editor. We report on an exploratory field study in which we found that participants were interested in whether code is recent but do not care precisely how recent and that a number of code characteristics lead to developer precaution, but change size is not one of them.

Title: How Should We Read and Analyze Bug Reports: An Interactive Visualization using Extractive Summaries and Topic Evolution
Author: Shamima Yeasmin, Chanchal K. Roy and Kevin A. Schneider

Abstract for How Should We Read and Analyze Bug Reports: An Interactive Visualization using Extractive Summaries and Topic Evolution
Software projects evolve over time as bugs are addressed and new functionalities are added. Managing bugs can be a significant challenge for a project manager especially when the number of reported bugs is large, and the manager needs to consult with them. It is also preferable that developers new to a project first familiarize themselves with the project and the reported bugs before actually working on them. In order to reduce developers' time and efforts for reading a bug report, in this paper, we propose a visualization technique that provides an extractive summary visualization for a given bug report. In addition, our proposed technique assists the developers or managers in reviewing a project's bug reports by interactively visualizing insightful information using topic analysis on the bug reports. In order to validate the effectiveness of our proposed visualization technique, we conducted a task-oriented user study involving six participants and a case study using 3914 bug reports. The findings from both studies show that our visualization technique is promising, and it can assist the comprehension and analysis of bug reports. The results from the user study indicate that visualized summary is relatively preferred to the non-visualized summary for quick comprehension of bug reports.

SMARTER SOFTWARE SYSTEMS/ Tuesday, November 3rd 11:15AM - 12:15PM (Jasmine)
Session chair: Mariano Consens

Title: Multitenancy Benefits in Application Servers
Author: Panagiotis Patros, Dayal Dilli, Kenneth Kent, Michael Dawson and Thomas Watson

Abstract for Multitenancy Benefits in Application Servers
Multitenancy enables sharing of resources between different users, also known as tenants and is a backbone feature of cloud computing. The tenants execute their code as if resources were held individually by them. The sharing is transparent; the tenants are isolated from each other and one tenant is not allowed to affect the performance of the rest by overusing a resource. We propose a theoretical model to describe and predict memory footprint reductions by different levels of multitenancy in application servers, including our multitenancy level, which enables even further sharing, acting as an Application-Server-as-a-Service (ASaaS). We confirm our model by implementing a small custom application server in Java and measuring its footprint for different multitenancy levels. We and that our ASaaS approach required up to 65% less memory without any major response time overheads. Finally, we perform an analysis of potential memory sharing on an enterprise software stack between different levels of multitenancy, including our proposed ASaaS level, and the results support our findings.

Title: Data-dependence Profiling to Enable Safe Thread Level Speculation
Author: Arnamoy Bhattacharyya, Jose Nelson Amaral and Hal Finkel

Abstract for Data-dependence Profiling to Enable Safe Thread Level Speculation
Data-dependence profiling is a technique that enables a compiler to judiciously decide when the execution of a loop | which the compiler could not prove to be dependence free | should be speculated through the use of Thread Level Speculation (TLS). The data collected by a data-dependence profiler can be used to predict if may dependencies reported by a compiler static analysis are likely to materialize at runtime. A cost analysis can then be used to decide that some loops with a lower probability of dependence should be speculatively parallelized. This paper addresses the question as to whether a loops' dependence behaviour changes when the input to the program changes | a study of 57 different benchmarks indicates that it usually does not change. Then the paper describes SpecEval, an automatic speculative parallelization framework that uses single-input data-dependence profiles to find speculation candidates in the SPEC2006 and PolyBench/C benchmarks. This paper also presents a performance evaluation of TLS implementation in IBM's Blue-Gene/Q supercomputer and shows that the performance of TLS is affected by several factors, including the number of speculated loops, the execution-time coverage of speculated loops, the miss-speculation overhead, the L1 cache miss rate and the effect on dynamic instruction path length.

Title: MOTL: a Textual Language for Trace Specification of State Machines and Associations
Author: Hamoud Aljamaan and Timothy Lethbridge

Abstract for MOTL: a Textual Language for Trace Specification of State Machines and Associations
In a model-driven development (MDD) environment where most or all of the source code is generated from models, there is a lack of tools for model-level tracing, in which generated execution traces can be linked to model level constructs. This lack of tools might inhibit the use of MDD, since it forces modelers to do dynamic analysis at a lower level of abstraction. In this paper, we propose a solution allowing modelers to textually specify the tracing of modeling constructs: attributes, state machines and associations using Umple. The resulting execution traces then contain model construct links. We describe the syntax and semantics of the language, as well as the generated execution traces, and give an example.

Theme

"CASCON 2015 Conference Theme: Redefining the New Enterprise"

Transformational technologies in cloud, mobile, cognitive computing, internet of things, social computing, analytics and big data and others have become vortexes that change the enterprise landscape. For enterprises to stay relevant

and competitive, facing the challenge of understanding and leveraging these new technologies is a must. Enterprises who can leverage new technologies to lead with new enterprise paradigms will set themselves apart from their competitors. They can deliver distinguishing appeals as they engage their customers. They can take the business to a newer level of success that others can't. Commerce, healthcare, finance, smarter cities are examples of industrial domains going through such technological evolution. This conference provides a forum to ideate, discuss, evaluate the depth of each transformational technology, as well as their confluence impact in redefining new enterprises.

Research Topics

Please expand below to see a descriptions of the research topic

Integrated Solution Areas and Applications
Smarter city solutions, industrial applications, healthcare solutions, media services

Smarter Systems and Services
Smart interactions, context-awareness, personal agents, end-user programming, mash-ups, web interaction frameworks, semantic web, ontologies

Internet of Things
Wearable computing, anticipatory computing, quantified self, gamification, personal analytics, smart devices, integration and connectivity, machine-to-machine platforms, embedded sensors, monitoring and control, security

Software as a Service
Systematic business and computing models, large-scale system integration, economic and societal issues, legal and privacy concerns

Mobile and Ubiquitous Computing
Mobile application platforms, mobile web, write once run anywhere, mobile application security, mobile application management

Social Computing
Collaboration environments, social analytics, interaction patterns, human-computer interaction

Cloud Computing and Infrastructure
Cloud platforms and frameworks, self-managed systems, adaptive systems, virtualization

Big Data and Databases of the Future
Big data, data integration, data warehouses, integrated data management, database technology, text data, stream data, XML, linked data, query processing and optimization

Business Intelligence and Analytics
Predictive analytics, text analytics, risk analytics, intelligent search, data mining, knowledge discovery, information visualization, OLAP

Software Processes and Tools

Agile methods, collaboration and awareness, open source, compilers and interpreters, scripting languages, development environments, service-oriented programming, distributed systems, ultra-large scale software systems

Software analysis and Modelling
Software analysis, mining software repositories, program comprehension, software testing, model checking, metrics

Dynamic Business Processes
Business process modeling, service-oriented architecture, REST services

Energy Aware Computing
Energy-aware platform architecture, energy-aware development tools, energy-aware resource management

Next Generation Systems
Multi-core systems, stream processing systems, graphical processing systems, FPGA

FAQ

Is it necessary to be a student to submit a paper? I will be graduating before the submission deadline.
You are eligible to submit. You do not need to be a student at the time.

What is an abstract?
This is a brief summary of the paper to allow readers to understand what the paper is about. The abstract is to be submitted through EasyChair on the web in a text field. You do not need to upload the abstract (or extended abstract) in PDF form. Submitted abstract entries are used to help assign reviewers.

Can I submit abstracts and technical papers at the same time before the abstract due date?
Yes.

What does technical papers notification mean?
Relevant papers to the conference are reviewed by at least three members of the technical program committee. Notification is telling the author(s) whether their paper has been accepted or rejected for publication after that process.

What are camera-ready technical papers?
This is a traditional term for the final version of an accepted paper, ready for publication. The paper must conform to the required formatting, and typically also addresses issues identified by the reviewers.

Is there a fee to register papers?
There is no fee for papers that had been accepted after the review process to be published in the conference proceedings or in the ACM Digital Library.

Can papers be published in absentia?
At least one author of an accepted publication must attend the conference in person to present the work.

How are papers to be presented?

If accepted, the work is to be presented in person at the conference event, typically with your own notebook computer in connection to a provided projector. Presenters typically have 20 minutes for both their presentation and audience questions.

Will the conference proceedings be published?

The proceedings (i.e., technical papers, position papers, and workshop descriptions) will be published digitally in the ACM Digital Library soon after the event.

What is the registration fee for the conference?

There is a registration sign-up process for the conference, but there is no actual registration fee for any attendee to pay.

IBM

Emerging Technologies Track (ETT)

ETT Call for papers, click here for more details

CASCON Emerging Technology Track invites position papers to share visions, ideas, and insights on these advancements, highlight technology gaps, research agendas, and integrated solutions, and initiate discussion, feedback, and further collaboration.

ETT Program Co-Chairs:

Ying (Jenny) Zou, Queen's University, Canada

Ken Wong, University of Alberta, Canada

ETT Agenda

ETT Agenda (117KB)

CASCON 2015 Accepted ETT Position Paper

Congratulations to the following authors on their ETT Position Paper being accepted by CASCON 2015.

Title: To Default or not to Default: Exposing Limitations to HBase Cluster Deployers
Authors: Roni Sandel, Marios Fokaefs, Mark Shtern, Hamzeh Khazaei and Marin Litoiu
Keywords: Cloud Computing, Big Data, Design, Performance Modelling and Evaluation, Migration

Abstract for To Default or not to Default: Exposing Limitations to HBase Cluster Deployers
With the advent of sensor networks and portable devices, data has been produced rapidly and in great amount. As a result storing and processing Big Data, in combination with the advances in cloud and virtual infrastructures, pose interesting challenges. In our previous work, we studied these challenges with various experiments around different HBase cluster configurations and their impact on the performance of the cluster. A by-product of our experiments was that, in spite of advances in tooling support to set up and configure a Big

Recent News

(NEW!) CASCON 2015 Workshop Materials are now available for Download!

(NEW!) CASCON 2015 Survey is now Available

(NEW!) CASCON 2016, Nov. 8th-10th. Mark you calendar. Program co-chairs for CASCON 2016 are Hausi A. Müller, University of Victoria and Vio Onut, IBM Canada CAS

(NEW!) Recorded keynotes, FoSP sessions and workshop presentations will be available for download after the conference. Stay tuned.

Reminder: CASCON 2015 Opening Keynote Session starts at 9AM on Monday, Nov 2nd, 2015

CASCON 2015 Program Guide is available for download Now

CASCON 2015 Proceedings is available for download Now

Need help on CASCON website? Contact cwillis@ca.ibm.com or casresearch@ca.ibm.com

Follow us through these social outlets:

Facebook

Google+

LinkedIn

Twitter

Related Links

The Four Languages of Innovation: A Model of Industrial Research and Development (1MB)

Research Partner

Data cluster, the various tools are not always aligned to produce the optimal or near-optimal performance for data clusters. More specifically, we show that the default configuration values of state-of-the-art cluster deployers, including Cloudera, IBM BigInsights, Apache Hortonworks and the manual HBase deployment, do not take in to account the underlying infrastructure resulting in subpar performance.

Title: Towards Improved Performance and Compliance in Healthcare using Wearables and Bluetooth Technologiess
Authors: Omar Badreddin, Ricardo Castillo, Lysanne Lessard and Michael Albanese
Keywords: Wearables, Healthcare Informatics, Bluetooth, Compliance and Governance, Indoor Localization

Abstract for Towards Improved Performance and Compliance in Healthcare using Wearables and Bluetooth Technologies
Performance and compliance in hospitals are particularly challenging due to the nature of care variability, and the human-intensive clinical activities involved. There is a need to collect fine-grained measurements to enact performance and governance effectively. Modern wearable technologies, such as smart watches, provide a significant opportunity towards facilitating the gathering of such fine-grained data points such as location. However, indoor localization using wearables is limited due to many factors, including signal sensitivity, interference, and variations in manufacturers specifications. This paper introduces a domain problem from healthcare, demonstrates empirically the limitations with today's wearables, and outlines novel approaches to dealing with such limitations. Specifically, we demonstrate how modern wearables can be deployed to collect fine-grained performance measures and facilitate governance.

Title: A Tangible User Interface for Interactive Data Analysis
Authors: Ana Jofre, Steve Szegeti, Stephen Tiefenbach Keller, Frederico Tome Filho, Lan-Xi Dong and Sara Diamond
Keywords: Tangible User Interface, Graspable User Interface, 3d Interactions, Interface Design, Database query, Data Analytics, Human Computer Interaction, Collaboration, Collaborative Learning, Collaborative Work Environment, Data Materialization

Abstract for A Tangible User Interface for Interactive Data Analysis
We present a prototype for a Tangible User Interface (TUI) designed to interactively query a database. While much work has been done on TUI, showing that they encourage collaboration and positively enhance user experience, few tangible systems have been designed specifically for data analysis tasks. Our system combines a tabletop user interface with a two-dimensional screen display; the user interrogates the data by placing tokens on or off the tabletop and the screen displays the results of the user's query. The objects are tagged and identified using fiducial markers that are read with open-source Reactivision computer vision software, and the visualization code is written in Processing. We use radio station listenership demographic data for this prototype, but the system can be used to query any type of demographic database. While data analytics is typically done alone, this tool encourages data exploration to be a group/team activity. We believe that collaborative work is beneficial because working together may be helpful in reducing miscommunications and misunderstandings. Since humans are naturally social animals, a collaborative work environment may lead to higher levels of engagement, thereby stimulating creativity. Greater engagement and creativity may also result from the playful nature of the tangible interaction.

Title: Social Computing and Intelligence: Opportunities for the Public and the Enterprise
Authors: Taraneh Khazaei, Lu Xiao, Robert Mercer and Atif Khan
Keywords: Social Computing, Collective Intelligence, Bussiness Intelligence, Persuasive Discourse, Social Privacy

Abstract for Social Computing and Intelligence: Opportunities for the Public and the Enterprise

Online social media have enabled mass data exchange at a global scale, leading to new and advanced forms of intelligence that arise from online content generation and interaction. In the form of collective intelligence, crowds gather together and contribute to solving problems that may be difficult or impossible to solve by individuals and single computers. In addition, the consumer insight revealed from the intelligence can be leveraged to build powerful business intelligence tools, enabling efficient and effective business decision-making processes. In this paper, we discuss the current research gaps, our research efforts to address these gaps, as well as suggested directions in social media mining and analytics to support collective and business intelligence.

Title: Evolutionary Analysis of Access-Control Models: A Formal Concept Analysis Method
Authors: Zhuobing Han, Mathieu Merineau, Francois Gauthier, Ettore Merlo, Xiaohong Li and Eleni Stroulia
Keywords: Software Evolution, Role Based Access Control (RBAC), Formal Concept Analysis (FCA)

Abstract for Evolutionary Analysis of Access-Control Models: A Formal Concept Analysis Method
Access control is an essential feature of most software-systems security mechanisms. Role-Based Access Control (RBAC), likely the most popular access-control technique, specifies "user roles" and associates each role with "permissions" to access distinct system functionalities. These role-permission assignment rules, as well as the types of system users and system functionalities, evolve over time. In this paper, we discuss a methdology for analyzing and understanding the RBAC-evolution process, its relation to the overall evolutionary lifecycle of the system, and its impact to any security vulnerabilities that the system may suffer from. Our methodology relies on Formal Concept Analysis (FCA). First, we extract the role-permissions matrix of each system version and we compute the implicit concept lattice. Next, we apply a suite of distance metrics to pairwise compare the matrix and concept-lattice versions. By examining the evolution of these distance metrics, developers can easily notice which versions involve more, and more complex, RBAC changes that may possibly indicate higher security risks. We demonstrate our methdoology with a study of the Mediawiki platform.

Title: Towards A Context-Aware Mobile App Management Framework
Authors: Ernest Aaron and Hausi Muller
Keywords: Mobile App Management, Smart Applications, App Recommender, Context-Aware Apps

Abstract for Towards A Context-Aware Mobile App Management Framework
In an app-driven society, the apps installed on smartphones inevitably will increase with time. As a result, navigating to locate apps, requires scrolling back and forth between various pages. Also, finding apps in the app store is complex due to the vast amount of choices. There is a need for mobile application management. Our proposed app management approach will mine contextual data from the user, information about apps installed, app usage, location and time. The data collected will be used to organize relevant location based apps dynamically in a Context-Aware App widget. Our framework, consist of a two tier app-recommender component, which provides users with additional apps that may be needed but are not installed. A coarse-grained app recommender filters nearby vendor services with apps not installed whereas the fine-grained app recommender, filters apps based on users' preferences in their Personal Context Sphere (PCS).

Title: Github's Big Data Adaptor: An Eclipse Plugin
Authors: Ali Sajedi Badashian, Vraj Shah and Eleni Stroulia
Keywords: GitHub, Mining software repositories, Eclipse Plugin, Software tools, Big data

Abstract for Github's Big Data Adaptor: An Eclipse Plugin
The data of GitHub, the most popular code-sharing platform, fits the characteristics of "big data" (Volume, Variety and Velocity). To facilitate studies on this huge GitHub data volume, the GHTorrent web-site publishes a MYSQL dump of (some) GitHub data quarterly.

Unfortunately, developers using these published data dumps face challenges with respect to the time required to parse and ingest the data, the space required to store it, and the latency of their queries. To help address these challenges, we developed a data adaptor as an Eclipse plugin, which efficiently handles this dump. The plugin offers an interactive interface through which users can explore and select any field in any table. After extracting the data selected by the user, the parser exports it in easy-to-use spreadsheets. We hope that using this plugin will facilitate further studies on the GitHub data as a whole.

Title: Harnessing the Power in Your Pocket
Authors: Joseph Spitzer, Kate Lockwood and Jason Sawin
Keywords: Smartphones, Mobile Devices, Grid Computing

Abstract for Harnessing the Power in Your Pocket
In the mid-90s when personal computers became common, organizations began to consider ways to harness the power of underutilized processor cycles. One of the most prominent of these efforts, SETI@home, recruited volunteers to donate their unused processor cycles to analyze radio telescope data, work that would have previously been done on special-purpose clusters. Now, individuals carry smartphones with processing capabilities that far outstrip those of a mid-90s desktop. At the same time, the amount of business data being collected is growing exponentially and corporations are scrambling to find the computing resources necessary to process it. In this paper we explore the feasibility of employers utilizing unused cycles on employee company-issued mobile devices, much like SETI@home did with personal computers, to process business intelligence data.

Title: Emerging technologies for Enterprise Linux on IBM z Systems
Authors: Bryan Chan, Rishi Misra and Lei Zhang
Keywords: IBM z Systems, Open source software ecosystem, Linux

Abstract for Emerging technologies for Enterprise Linux on IBM z Systems
The advent of mobile and cloud computing in the past decade has led to an explosion of innovative open source technologies. Some examples are NoSQL databases such as Apache Cassandra, Apache CouchDB and MongoDB, message brokers such as RabbitMQ and Apache ActiveMQ, new programming languages and run-times such as Go and server-side JavaScript, and IT automation and cloud infrastructure tools such as Puppet, Docker and Cloud Foundry. As enterprise Linux users start to embrace these emerging technologies, they are discovering that many open-source developers and vendors only support x86-based platforms, and not enterprise server platforms such as those based on IBM z Systems and IBM Power Systems. In this position paper, we discuss some of the challenges that we have encountered while closing the gap in emerging technologies between IBM z Systems and x86, present some early results from our effort, and indicate future directions for our work.

Title: Context-aware Real-time Video Analytics
Authors: Nina Taherimakhsousi and Hausi A. Muller
Keywords: Context-based Video Search, Context-awareness, Decentralized, Cloud Computing, Real-time Video Analytics

Abstract for Context-aware Real-time Video Analytics
Cameras are found everywhere in the Internet of Things (IoT). We propose a scalable platform using SAVI for the continuous collection of context-aware processed video from mobile devices. Users can perform efficient context-based searches on the total collection of labeled videos. In this paper we discuss the challenges and benefits for video upload, video labeling and context-based search on the SAVI network.

Title: Pilot Study of Collective Decision-making in Code Review Process
Authors: Toshiki Hirao, Akinori Ihara and Kenichi Matsumoto
Keywords: Collective Decision-making, Code Review Process, Simple Majority Method

Abstract for Pilot Study of Collective Decision-making in Code Review Process
Political argument often follow the simple majority method which is a decision rule that selects alternatives which have a majority, that is, more than half the votes. On the other hand,

Today's Q¥&A services such as StackExchange and Yahoo! Answers are working the simple majority method only as reference. For example, if an answer gets majority voting (which means a favorite answer for watchers), the questioner might like another answer. Then, the voting would be used as reference. Open Source Software (OSS) projects also work the voting approach only as reference from developers (reviewers) to verify a changed source code (patch). However, the OSS projects might not follow the majority voting. This study identifies how many code reviews based on collective decision-making follow the simple majority method. We have conducted a case study using Qt project data to analyze the criteria to integrate a patch in OSS projects. From the result, we found that the patch with negative votes is likely to be rejected. And, latter voting in reviewers' discussion would be impact to the final decision.

IBM

Workshops

Workshops are held throughout the conference in the form of tutorials and hands-on. Workshop hosts come from a variety of backgrounds. Any individual is free to submit an application and request to host a workshop. The workshops are then selected and the public is given notice of the workshops that will be available during the conference. Early registration is required as seating is limited to a first-come-first-served basis. Registration to the conference and its workshops are available on the CASCON site.

CASCON 2015 Workshop Materials

CASCON 2015 Workshop Materials(PDF)

CASCON 2015 Accepted Workshop

Congratulations to the following chairs on their Workshops being accepted by CASCON 2015.

Workshop Schedule : Monday, November 2nd

Morning 10:45AM - 1:00PM
Title: IBM Community Hackathon 2015 - Bluemix (IBMers-Only)
Location: Conf Center 1
Theme: Cloud Computing
Chairs: Vince Yuen, Chinh Hoang, Allen Chan, Lucy Li, Riley Draward

Afternoon 2:00PM - 5:00PM
Title: Workshop on Financial Risk Analytics
Location: Conf Center 1
Theme: Data and Analytics
Chairs: Oleksandr Romanko, Alex Kreinin, Helmut Mausser

Title: Remote Patient Monitoring to Improve Health: Challenges and Opportunities
Location: Willow 3
Theme: Data and Analytics
Chairs: Randy Giffen, Michael Fung-Kee-Fung, Mihaela Rotaru

Recent News

(NEW!) CASCON 2015 Workshop Materials are now available for Download!

(NEW!) CASCON 2015 Survey is now Available

(NEW!) CASCON 2016, Nov. 8th-10th. Mark you calendar. Program co-chairs for CASCON 2016 are Hausi A. Müller, University of Victoria and Vio Onut, IBM Canada CAS

(NEW!) Recorded keynotes, FoSP sessions and workshop presentations will be available for download after the conference. Stay tuned.

Reminder: CASCON 2015 Opening Keynote Session starts at 9AM on Monday, Nov 2nd, 2015

CASCON 2015 Program Guide is available for download Now

CASCON 2015 Proceedings is available for download Now

Need help on CASCON website? Contact cwillis@ca.ibm.com or casresearch@ca.ibm.com

Follow us through these social outlets:

Facebook

Google+

LinkedIn

Twitter

Related Links

The Four Languages of Innovation: A Model of Industrial Research and Development (1MB)

Research Partner

Title: Agile Software Development for Bluemix with IBM DevOps Services

Location: Holly + Butternut

Theme: Cloud Computing

Chairs: Kris Kobylinski, Ian He, Nancy A. Schipon

Title: Hands-on: Introduction to Developing Java EE 7 Applications using WebSphereÂ® Application Server Developer Tools for Eclipse (WDT) and Liberty Profile

Location: Orchid

Theme: Next Generation Systems

Chairs: Elson Yuen, Erin Harris, Keith Chong, Vaninder Rajput, Rajiv Senthilnathan

---

Title: Building Your Own Language Runtime

Location: Jasmine

Theme: Next Generation Systems

Chairs: Angela Lin

Title: Data Science Workshop: Experience Driven Analytics

Location: Evergreen

Theme: Other

Chairs: Joanna Ng, Frank Dehne, Stan Matwin, Herna Viktor, Olga Baysal

---

Title: Role play - Learn to secure resources and APIs effectively using OAuth and OpenID Connect standards

Location: Primrose

Theme: Security Platform and Tools

Chairs: Krithika Prakash

Title: Fourth Annual Mobile Application Development with IBM Bluemix and IBM Mobile First Platform (Repeating)

Location: Violet

Theme: Smart Interaction Ecosystems

Chairs: Salman Moghal, Serjik G. Dikaleh, Ozair Sheikh

Workshop Schedule : Tuesday, November 3rd

Afternoon 2:00PM - 5:00PM

Title: 10th Workshop on Challenges for Parallel Computing

Location: Conf Center 1

Theme: Next Generation Systems

Chairs: Robert Ho, Jeremy Bradbury, Kit Barton

Title: Workshop on the Application of Security and Testing to Rich Internet Applications

Location: Willow 3

Theme: Security Platform and Tools

Chairs: Guy-Vincent Jourdan, Gregor Bochmann, Ettore Merlo, James Miller, Vio Onut, Lin Tan

---

Title: Introduction to debugging and monitoring Node.js

Location: Holly + Butternut

Theme: Cloud Computing

Chairs:Gary Liu, Joran Siu, Michael Dawson, Ivy Ho, Junliang Yan

---

Title: Hands-on: Developing Java and JavaScript Applications on Bluemix using IBM Eclipse Tools for Bluemix

Ontario Centres of Excellence (OCE)

Consortium for Software Engineering Research (CSER)

National Research Council Canada (NRC)

Interested in Becoming a Faculty Fellow?

Click Here to Read about the Requirements

.

Location: Orchid

Theme: Cloud Computing

Chairs: Jonathan West, Ben Chen, Steven Hung, Elson Yuen

Title: DB2 LUW Native Encryption

Location: Violet

Theme: Security Platform and Tools

Chairs: Mihai Nicolae, Mihai Iacob

Title: Developing next generation analytics applications on LinuxOne end-to-end Architecture and Implementation

Location: Jasmine

Theme: Next Generation Systems

Chairs: Elton de Souza

Title: Collecting and analyzing depression notes using IBM Social Media Analytics

Location: Evergreen

Theme: Data and Analytics

Chairs: Diman Ghazi, Yvon Leclerc

Title: Commercialization of Technology Research for Benefit

Location: Primrose

Theme: Other

Chairs: Kathryn Brohman, Paul Ward

Workshop Schedule : Wednesday, November 4th

Morning 10:00AM - 1:00PM
Title: 14th Compiler-Driven Performance Workshop

Location: Conf Center 1

Theme: Next Generation Systems

Chairs: Ondrej Lhotak, Xipeng Shen, Ettore Tiotto, Martin Hirzel, Clark Verbrugge

---

Title: Introducing IBM Watson Analytics into Curricula and Programs

Location: Primrose

Theme: Data and Analytics

Chairs: Dennis Buttera, Jennifer Collins

Title: The 7th CASCON Workshop on Cloud Computing

Location: Conf Center 2

Theme: Cloud Computing

Chairs: Marin Litoiu, Joe Wigglesworth

Title: Getting Started with DB2 V10.5 BLU

Location: Jasmine

Theme: Data and Analytics

Chairs: David Dohyung kim, Anthony Reina

Title: Fourth Annual Mobile Application Development with IBM Bluemix and IBM Mobile First Platform

Location: Holly + Butternut

Theme:Smart Interaction Ecosystems

Chairs: Salman Moghal, Serjik G. Dikaleh, Ozair Sheikh

Title: Business Process Management in a Day

Location: Orchid

Theme: Smart Interaction Ecosystems

Chairs: Gary Bist

Title: Taming a Tiger - Software Engineering in the Era of Big Data & Continuous Development

Location: Evergreen

Theme: Next Generation Systems

Chairs: Craig Statchuk, Nazim Madhavji, Andriy Miranskyy, Frank Dehne

Afternoon 1:00PM - 2:00PM
Title: WiT - Memoirs from a Nobody

Location: Conf Center 1

Theme: Other

Chairs: Maria Attarian, Paul Zikopoulos, Joanna Ng

Afternoon 2:00PM - 5:00PM
Title: 14th Compiler-Driven Performance Workshop

Location: Conf Center 1

Theme: Next Generation Systems

Chairs: Ondej LhotÃ¡k, Xipeng Shen, Ettore Tiotto, Martin Hirzel, Clark Verbrugge

---

Title: Introducing IBM Watson Analytics into Curricula and Programs

Location: Primrose

Theme: Data and Analytics

Chairs: Dennis Buttera, Jennifer Collins

Title: Engineering Cyber Physical Systems

Location: Conf Center 2

Theme: Smart Interaction Ecosystems

Chairs: Hausi Muller, John Mylopoulos, Marin Litoiu

Title: Building IoT-enabled applications

Location: Holly + Butternut

Theme: Other

Chairs: Victor Sosa, Orlando Rincon, Cesar Orozco

Title: Business Process Management in a Day

Location: Orchid

Theme: Smart Interaction Ecosystems

Chairs: Gary Bist

Title: Taming a Tiger - Software Engineering in the Era of Big Data & Continuous Development

Location: Evergreen

Theme: Next Generation Systems

Chairs: Craig Statchuk, Nazim Madhavji, Andriy Miranskyy, Frank Dehne

---

Title: Hands-on: Developing Java and JavaScript Applications on Bluemix using IBM Eclipse Tools for Bluemix (Repeating)

Location: Violet

Theme: Cloud Computing

Chairs: Jonathan West, Ben Chen, Steven Hung, Elson Yuen

IBM

Technology Showcase

The Technology Showcase is open throughout the 3-day conference and includes poster displays and demonstrations of new technological innovations. These displays demonstrate on-going research projects of graduate students from universities as well as those from industry. It allows communication and networking in a non-formal setting.

Technology Showcase Map

Technology Showcase Map (554KB)

CASCON 2015 Accepted Technology Showcase

Congratulations to the following exhibitors/co-exhibitors on their Technology Showcases being accepted by CASCON 2015.

Data and Analytics

Title: Gestalt Laws of Grouping based Visual Similarity and Web Page Classification Model
Exhibitors/Co-Exhibitors: Zhen Xu, Syed Tauhid Zuhori, James Miller

Abstract for Gestalt Laws of Grouping based Visual Similarity and Web Page Classification Model
Traditional text-based web page classification fails to handle rich-information-embedded modern web pages. This paper proposes a methodology to evaluate modern web page similarity by visual information using Gestalt laws of grouping, and to classify them in terms of their visual similarity.

Title: Gestalt Laws of Grouping based Web Page Semantic Block Identification
Exhibitors/Co-Exhibitors: Zhen Xu, Fadwa Estuka, James Miller

Abstract for Gestalt Laws of Grouping based Web Page Semantic Block Identification
Semantic block identification is an approach to retrieve information from web pages and applications. As website design evolves, however, traditional methodologies cannot perform well any more. This paper proposes a new model to merge web page content into semantic blocks by simulating human perception.

Title: Monitoring business goals with Argus
Exhibitors/Co-Exhibitors: Alejandro Mate

Recent News

(NEW!) CASCON 2015 Workshop Materials are now available for Download!

(NEW!) CASCON 2015 Survey is now Available

(NEW!) CASCON 2016, Nov. 8th-10th. Mark you calendar. Program co-chairs for CASCON 2016 are Hausi A. Müller, University of Victoria and Vio Onut, IBM Canada CAS

(NEW!) Recorded keynotes, FoSP sessions and workshop presentations will be available for download after the conference. Stay tuned.

Reminder: CASCON 2015 Opening Keynote Session starts at 9AM on Monday, Nov 2nd, 2015

CASCON 2015 Program Guide is available for download Now

CASCON 2015 Proceedings is available for download Now

Need help on CASCON website? Contact cwillis@ca.ibm.com or casresearch@ca.ibm.com

Follow us through these social outlets:

Facebook

Google+

LinkedIn

Twitter

Related Links

The Four Languages of Innovation: A Model of Industrial Research and Development (1MB)

Research Partner

Abstract for Monitoring business goals with Argus
Argus is a data analytics tool that allows users to monitor business goals through Key Performance Indicators. Argus exploits the multidimensionality of data to monitor data subspaces, thereby enabling users to detect problems that would otherwise be overlooked in the aggregates. It works both on static and streaming data.

Title: Queen's Drill Analytics
Exhibitors/Co-Exhibitors: Shadi Khalifa

Abstract for Queen's Drill Analytics
In this Demo, we introduce the Queen's Drill Analytics, a modified Apache Drill that allows calling WEKA analytics algorithms from within the Drill SQL statements and seamlessly distribute these algorithms over the Drill nodes.

Title: User Context Aware Service Composition
Exhibitors/Co-Exhibitors: Yu Zhao, Shaohua Wang, Ying Zou, Joanna Ng, Tinny Ng

Abstract for User Context Aware Service Composition
To select services satisfying user's preferences, we propose an approach to automatically extract user's goal and constraints from goal description, identify personalized web task collection from user's goal, and compose services to meet user's multiple preferences.

Title: Evaluating Cluster Configurations for Big Data Processing: An Exploratory Study
Exhibitors/Co-Exhibitors: Marios Fokaefs

Abstract for Evaluating Cluster Configurations for Big Data Processing: An Exploratory Study
We present a series of experiments to evaluate the impact of BigData cluster configurations to the cluster's performance. Our experiments on HBase clusters have shown that the topology of the cluster and the choice of data schema significantly impact the cluster's performance, which is linear to the workload.

Title: Collecting and analyzing depression notes using IBM Social Media Analytics
Exhibitors/Co-Exhibitors: Diman Ghazi

Abstract for Collecting and analyzing depression notes using IBM Social Media Analytics
We explore the potential of using IBM Social Media Analytics to collect and analyze distress notes. Due to the billions of social media posting every day, an ontology will be defined to collect and curate only relevant notes. Next, we apply text analysis to define possible factors contributing to depression.

Title: Understanding Challenges Towards Generalizing Defect Prediction Models
Exhibitors/Co-Exhibitors: Feng Zhang, Ying Zou

Abstract for Understanding Challenges Towards Generalizing Defect Prediction Models
Most companies do not have sufficient resources to build and apply a defect prediction model. One major barrier to adopt defect prediction models in industry is the low generalizability of defect prediction models. Generalizing defect prediction models can save the effort to rebuild prediction models for a particular project.

Title: Entity Resolution with Graphs
Exhibitors/Co-Exhibitors: Lena Woolf , Dmitry Drinfeld, Andrew Wang, Christine Li

Abstract for Entity Resolution with Graphs
In this exhibit we discuss benefits of using graph databases and how performing entity

Ontario Centres of Excellence (OCE)

Consortium for Software Engineering Research (CSER)

National Research Council Canada (NRC)

Interested in Becoming a Faculty Fellow?

Click Here to Read about the Requirements

.

resolution on graph data can help customers discover new analytical insight

Cloud Computing

Title:Power Systems Academic Initiative

Exhibitors/Co-Exhibitors: Madeeha Ali, Nihal Perera, Bruce Turgeon

Abstract for Power Systems Academic Initiative
I would like to be able to allow different schools understand this no-charge initiative, that will allow them both to enhance and create courses relevant to today's career demands. I will be showing them the no-charge courses developed by IBM. I would also like to do a demo - POWER8.

Title: Connecting your devices on Cloud

Exhibitors/Co-Exhibitors: Victor Sosa, Orlando RincÃ³n, Cesar Orozco

Abstract for Connecting your devices on Cloud
Internet of Things is a reality that is already happening in the world. Everything is becoming more interconnected and technologists now have the opportunity to make it intelligent with the appropriate tools.

Title: Elasticsearch, Logstash, and Kibana with Linux on IBM System Z

Exhibitors/Co-Exhibitors: Bill O'Farrell, Aditya Murray, Sam Ding

Abstract for Elasticsearch, Logstash, and Kibana with Linux on IBM System Z
This demo is to act as a proof, for the successful porting of the ELK stack, which consists of Elasticsearch, Logstash and Kibana, to the IBM z System. The ELK stack provides a very powerful way to combine, filter and analyze data, using open-source software on System Z.

Title: Engineering Economics of Everything-as-a-Service (XaaS)

Exhibitors/Co-Exhibitors: Marios Fokaefs

Abstract for Engineering Economics of Everything-as-a-Service (XaaS)
The flexibility, offered by virtual infrastructures, has given rise to new business models and interesting implications in the economics of software ecosystems. We present an economic model for web systems on the cloud and how it can be used to make decisions about real-time adaptation of cloud applications.

Title: A Container Orchestration Method to Enhance Failure Resiliency in cloud platforms

Exhibitors/Co-Exhibitors: SeyedAli Jokar

Abstract for A Container Orchestration Method to Enhance Failure Resiliency in cloud platforms
Everyday more applications are being moved to containers to provide independent and easily deployable micro-services on the cloud. In this work, we present a method to increase resiliency of containers to failures by considering underlying virtual or physical resources.

Title: Distributed Fault-Tolerant P2P RIA Crawling

Exhibitors/Co-Exhibitors: Khaled Ben Hafaiedh, Gregor von Bochmann, Guy-Vincent, Jourdan, Iosif Viorel Onut

Abstract for Distributed Fault-Tolerant P2P RIA Crawling
Rich Internet Applications (RIAs) have been widely used on the Web as they were found to be responsive and user friendly. Distributed RIA crawling has been introduced with the aim to decrease the crawling time. In this exhibit, we introduce an efficient distributed RIA crawling system that is fault tolerant

Title: Self-adaptation and lifecycle management with Openstack Heat

Exhibitors/Co-Exhibitors: Marios Fokaefs

Abstract for Self-adaptation and lifecycle management with Openstack Heat
Heat is the Openstack orchestration service. It enables the abstract definition of cloud

topologies and connects all the other Openstack services to deploy the topologies. Using monitoring and management services, Heat also enables the self-adaptation of deployments. We extend Heat to include proactive adaptation and adaptation on the application level.

Title: Emerging Trends in Object Storage
Exhibitors/Co-Exhibitors: Carlos Cavanna, Ao Wan

Abstract for Emerging Trends in Object Storage
Shingled magnetic recording (SMR) drives have the potential to reduce storage costs. They have different recording characteristics than standard drives. We will talk about device-managed and host-managed drives. Storlets enable running user-defined computational modules as data is transferred to/from object storage. We will show storlets integration with Swift Object Storage.

Title: Applying Control Synthesis to Web Service Composition
Exhibitors/Co-Exhibitors: Francis Atampore

Abstract for Applying Control Synthesis to Web Service Composition
The use of control theory for the synthesis of provably correct Web service compositions is investigated and initial result is reported. The audience will learn about discrete-event control theory, and its extension and application to automatic Web service composition. An E-business example will be used for illustrating this proposal.

Smart Interaction Ecosystems

Title: Learning to Reuse User Inputs in Service Composition
Exhibitors/Co-Exhibitors: Shaohua Wang

Abstract for Learning to Reuse User Inputs in Service Composition
To help users fill redundant information into services, we propose an efficient learning-to-rank approach which captures user's previous inputs, analyzes user's history and contexts, and learns to recommend values to input parameters during the service composition.

Title: Sensorian Shield: An IoT Plug-in Module for the Raspberry Pi
Exhibitors/Co-Exhibitors: Qusay Mahmoud, Dhimiter Qendri, Michael Lesicisin

Abstract for Sensorian Shield: An IoT Plug-in Module for the Raspberry Pi
The Sensorian Shield is a plug-in module that extends your Raspberry Pi and is designed to help you turn your ideas into reality without the use of a breadboard for prototyping. It supports Node-RED.

Next Generation Systems

Title: Kaleidoscope
Exhibitors/Co-Exhibitors: Andreas Bergen, Hausi MÃ¼ller

Abstract for Kaleidoscope
Increasing network traffic, especially from video services, requires high bandwidth capabilities. SDN and NFV allow us to improve throughput and user experience when consuming customised high quality, high bandwidth video over the internet. By using SDN and NFV we demonstrate improved network usage, throughput and user experience for video applications.

Title: The OpenPOWER Foundation: Enabling Data-Centric Transformation of Data Centers, Cloud Providers, and HPC Clients
Exhibitors/Co-Exhibitors: David K. Tam, Douglas Gibbs, Neil Graham

Abstract for The OpenPOWER Foundation - Enabling Data Centric Transformation for Data Centers, Cloud Providers and HPC Clients
The Foundation's open Power architecture gives 150+ members from 23 countries the ability to innovate utilizing the entire hardware-software stack. Year's Highlights (1) 136% increase, 88 new members, (2) Two next generation supercomputer wins worth $502M with US DOE

and UK Government, (3) Wins across data centers, CSPs and MSPs.

Title: SyncDebugger: Automatic bug localization in multithreaded programs
Exhibitors/Co-Exhibitors: Jeremy Bradbury, Jeremy Kwok

Abstract for SyncDebugger: Automatic bug localization in multithreaded programs
SyncDebugger is an automated tool for identifying the location of bugs in multithreaded Java programs. SyncDebugger works by systematically noising (inserting delays) into different parts of a program in order to isolate the source code that contains a bug. We evaluate SyncDebugger using programs from the IBM Concurrency Benchmark.

Security Platform and Tools

Title: A New Automatic Methodology for User Behavior Profiling of Website
Exhibitors/Co-Exhibitors: Syed Tauhid Zuhori, James Miller

Abstract for A New Automatic Methodology for User Behavior Profiling of Website
We present an automatic methodology that can mine User behavior profile using economically available web server logs without adding any overloads to existing system. We inspected top popular websites and introduced nine most abundant user behavior profiles and classify the user trace into a UBP based on HMM based classification.

Title: Incremental Parsing applied on IBM Security AppScan Source
Exhibitors/Co-Exhibitors: Mathieu Merineau

Abstract for Incremental Parsing applied on IBM Security AppScan Source
This project investigates fine grain incremental approaches to increase the performance of IBM Security AppScan Source. We conceived an incremental parsing algorithm and implemented it in a research prototype based on AppScan. We evaluated it on WebGoat and obtained parsing time improvements ranging from 22% to 82%.

Title: Automatically Identifying Login and Logout Activities on Websites
Exhibitors/Co-Exhibitors: Fadwa Estuka, James Miller

Abstract for Automatically Identifying Login and Logout Activities on Websites
We're interested in locating sections of a web application that can be used to login or logout from, depending on a number of visual clues - DOM tree, and dynamic clues like cookies; and in observing actions that happen on them for security testing, and surveying purposes of e-commerce websites.

Title: The Reconstruction of User-Sessions from HTTP Traces in RIAs
Exhibitors/Co-Exhibitors: Salman Hooshmand, Muhammad Faheem

Abstract for The Reconstruction of User-Sessions from HTTP Traces in RIAs
Given the previously recorded HTTP traces of a RIA (rich internet application) , we propose a system which reconstructs user-interactions including clicks and user-inputs. The process is useful in many ways especially for security tools.

Title: Grouping and Visualizing Security Findings
Exhibitors/Co-Exhibitors: Jinqiu Yang, Lin Tan

Abstract for Grouping and Visualizing Security Findings
We propose an automatic technique to group and visualize security findings from AppScanSource.

Title: OpenFlow-enabaled DDoS Attack Mitigation on Hybrid Clouds
Exhibitors/Co-Exhibitors: Nasim Beigi, Cornel Barna, Mark Shtern, Hamzeh Khazaei, Marin Litoiu, Tinny Ng

Abstract for OpenFlow-enabaled DDoS Attack Mitigation on Hybrid Clouds
We design and implement a DDoS attack mitigation solution that protects applications on a public cloud by leveraging a private cloud and SDN overlay networks. The mitigation solution isolates the suspicious traffic on demand from the regular application clusters through

dynamically redirecting the suspicious traffic to the private cloud.

Title: Manage APIs using OAuth and Social Login

Exhibitors/Co-Exhibitors: Krithika Prakash, Ozair Sheikh

Abstract for Manage APIs using OAuth and Social Login

APIs are taking over the world and the cloud also. Securing APIs is one of the major area of concerns. OAuth is the industry wide standard being used to manage authorization to APIs. Social Login allows the capability to sign on using Social Login providers like Google, Facebook, LinkedIn, Twitter.

Other Themes

Title: Context Awareness in Video Analytics

Exhibitors/Co-Exhibitors: Nina Taherimakhsousi, Hausi A. MÃ¼ller

Abstract for Context Awareness in Video Analytics

We propose a scalable platform for the continuous collection of context-aware processed video from mobile devices. Users can perform efficient context-based searches on the total collection of labeled videos. In this exhibit we discuss the benefits and challenges for video context-aware labeling and video context-based search on the SAVI network.

Title: Test Case Prioritization Using Extended Digraphs

Exhibitors/Co-Exhibitors: Seyedeh Sepideh Emam

Abstract for Test Case Prioritization Using Extended Digraphs

In this study we propose an extended digraph model as the basis of a new test case prioritization technique. The model prioritizes test cases based upon the amount of change they cause in applications' states. This technique is a novel and applicable approach in regression testing procedure.

Title: ReHMM: A Hybrid Approach to Inferring Extended FSA Models from Software Executions

Exhibitors/Co-Exhibitors: Seyedeh Sepideh Emam

Abstract for ReHMM: A Hybrid Approach to Inferring Extended FSA Models from Software Executions

In this study, we propose applying a new algorithm to infer EFSAs using a novel hybrid approach to infer an extended form of finite state machines. Deriving behavioural models from software executions is a common approach used in supporting a broad range of software development, maintenance, verification and validation tasks.

Title: Automated Testing of Motion-based Events in Mobile Application

Exhibitors/Co-Exhibitors: Seyedeh Sepideh Emam

Abstract for Automated Testing of Motion-based Events in Mobile Application

In this study we propose a novel approach to automatically generate movement-based gestures in mobile applications. A HMM classifier is used to generate movements, which mimic a user's behaviour in interacting with the application's user interface (UI). This approach reduces the probability of developing error- prone motion-based mobile applications.

Title: Public Library Recommendation

Exhibitors/Co-Exhibitors: Jason Morawski

Abstract for Public Library Recommendation

Rural public libraries have difficulty making use of recommender systems due to the sparsity of users and content causing a perpetual "cold-start" scenario. Greater recommender performance is achieved by applying a hybrid method utilizing both content-based recommendation and collaborative filtering.

Title: Novel Point of Interest Recommendation with Location Based Social

Networks

Exhibitors/Co-Exhibitors: Jason Morawski

Abstract for Novel Point of Interest Recommendation with Location Based Social Networks

Location-based social networks record user check-ins. By analyzing the social, temporal, and spatial components of check-ins recommendations can be made for where a user may like to go in the future. The system can be used to recommend novel points of interest, which the user has never visited before.

Title: Test Smarter, Test Early, Test Often

Exhibitors/Co-Exhibitors: Dorra Bouchiha, Nemanja Ivanovic, Ray Kivisto

Abstract for Test Smarter, Test Early, Test Often

Adequate testing is often a major stumbling block for DevOps adoption. We present a multi-tiered testing approach based on code coverage that provides targeted testing for each changeset, duplicate-free testing for daily builds and full testing for weekly builds. The cycle of catching and fixing regressions is thereby significantly shortened.

Title: IBM Vector Library

Exhibitors/Co-Exhibitors: Ian McIntosh, Tian Wang

Abstract for IBM Vector Library

To get maximum performance on Intel systems, many programs now use the SSE and AVX SIMD Vector functions. Unfortunately that makes the program hard to port elsewhere. One solution is the IBM Vector Library, providing SSE and AVX semantics on PowerPC AIX and Linux using the XLC or GCC compilers.

IBM

Conference Chair

Marcellus Mindel, Director of CAS, IBM Canada Lab

Joanna Ng, Director of CAS Research, IBM Canada Lab

Program Committee

Program Co-chairs

Hanan Lutfiyya, Western University, Canada

Marin Litoiu, York University, Canada

Steering Committee

Hans-Arno Jacobsen, University of Toronto, Canada

Hausi Müller, University of Victoria, Canada

Ying (Jenny) Zou, Queen's University, Canada

James R. Cordy, Queen's University, Canada

Jimmy Lo, CAS Research, IBM Canada Lab

Joanna Ng, CAS Research, IBM Canada Lab

Krzysztof Czarnecki, University of Waterloo, Canada

Marin Litoiu, York University, Canada

Hanan Lutfiyya, University of Western Ontario, Canada

Ken Wong , University of Alberta, Canada

Organization Team

Joanna Ng, CAS Research, IBM Canada Lab

Jimmy Lo, CAS Research, IBM Canada Lab

Tinny Ng (Workshop Chair), CAS Research, IBM Canada Lab

Criag Statchuk, CAS Research, IBM Canada Lab

Vio Onut, CAS Research, IBM Canada Lab

Debbie Kilbride, CAS Research, IBM Canada Lab

Recent News

(NEW!) CASCON 2015 Workshop Materials are now available for Download!

(NEW!) CASCON 2015 Survey is now Available

(NEW!) CASCON 2016, Nov. 8th-10th. Mark you calendar. Program co-chairs for CASCON 2016 are Hausi A. Müller, University of Victoria and Vio Onut, IBM Canada CAS

(NEW!) Recorded keynotes, FoSP sessions and workshop presentations will be available for download after the conference. Stay tuned.

Reminder: CASCON 2015 Opening Keynote Session starts at 9AM on Monday, Nov 2nd, 2015

CASCON 2015 Program Guide is available for download Now

CASCON 2015 Proceedings is available for download Now

Need help on CASCON website? Contact cwillis@ca.ibm.com or casresearch@ca.ibm.com

Follow us through these social outlets:

Facebook

Google+

LinkedIn

Twitter

Related Links

The Four Languages of Innovation: A Model of Industrial Research and Development (1MB)

Research Partner

Brenda Chow, CAS Research, IBM Canada Lab

Jordan Gould, CAS Research, IBM Canada Lab

Willis Chien, CAS Research, IBM Canada Lab

Xianming(Sam) Luo, CAS Research, IBM Canada Lab

Rami Matloob, CAS Research, IBM Canada Lab

Ivy Ho (Technology Showcase Chair), IBM Canada Lab

Sean Mulhern, IBM Canada Lab

Arif Saleh, IBM Canada Lab

Chad Corbitt, IBM Canada Lab

Alan Heighway, IBM Canada Lab

Rodney d'Silva, IBM Canada Lab

James Babarinde, IBM Canada Lab

Kaitlyn Pannunzio, IBM Canada Lab

Matina Kalkounis, Blackbox Communications

Program Committee Members

Marin Litoiu (PC Chair) , York University, Canada

Hanan Lutfiyya (PC Chair) , University of Waterloo, Canada

Alex Thomo , University of Victoria, Canada

Ashraf Aboulnaga , Qatar Computing Research Institute

Anwar Haque , University of Waterloo, Canada

Aleksander Essex , Western University, Canada

Calisto Zuzarte , IBM Canada Lab

Catia Trubiani , Gran Sasso Science Institute

Carl Worms , Credit Suisse AG, Information Technology

Cristiana Amza , University of Toronto, Canada

Daqing Hou , Clarkson University, United States

Dwight Makaroff , University of Saskatchewan, Canada

Dorina Petriu , Carleton University, Canada

Diwakar Krishnamurthy , University of Calgary, Canada

Ding Yuan , University of Toronto, Canada

Ebrahim Bagheri , Ryerson University, Canada

Eleni Stroulia , University of Alberta, Canada

Eyal De Lara , University of Toronto, Canada

Frank Dehne , Carleton University, Canada

Fred Popowich , Simon Fraser University, Canada

Grace Lewis , Carnegie Mellon Software Engineering Institute

Ontario Centres of Excellence (OCE)

Consortium for Software Engineering Research (CSER)

National Research Council Canada (NRC)

Interested in Becoming a Faculty Fellow?

Click Here to Read about the Requirements

.

Hausi A. Muller , University of Victoria, Canada

Ioana Baldini , IBM Research

Iosif Viorel Onut , IBM and University of Ottawa, Canada

Janice Singer , National Research Council

James R. Cordy , School of Computing, Queen's University, Canada

Jens Weber , University of Victoria, Canada

Joe Wigglesworth , IBM Canada Lab

Julia Rubin , Massachusetts Institute of Technology, United States

Jose Nelson Amaral , University of Alberta, Canada

Kenny Wong , University of Alberta, Canada

Kelly Lyons , University of Toronto, Canada

Ken Salem , University of Waterloo, Canada

Kostas Kontogiannis , U. Waterloo / NTUA, Canada

Ladan Tahvildari , University of Waterloo, Canada

Mariano Consens , University of Toronto, Canada

Massimiliano Di Penta , Dept. of Engineering - University of Sannio, Italy

Maurizio Morisio , Politecnico di Torino, Italy

Merve Unuvar , IBM Research

Mohammad Zulkernine , Queen's University, Canada

Mohammad Sadoghi , IBM T.J. Watson Research Center

Michael Smit , Dalhousie University, Canada

Mike Mcallister , Dalhousie University, Canada

Michael Bauer , Western University, Canada

Michael Kwok , IBM Canada

Norha M. Villegas , Icesi University, Cali, Colombia

Nick Cercone , York University, Canada

Nicholas Graham , Queen's University, Canada

Paul Ward , University of Waterloo, Canada

Pat Martin , Queen's University, Canada

Samuel Kounev , University of Wuerzburg, Germany

Sudhakar Ganti , University of Victoria, Canada

Tinny Ng , IBM Canada

Uwe Assmann , TU Dresden, Germany

Vassilios Tzerpos , York University, Canada

Wendy Powley , Queen's University, Canada

Yelena Yesha , University of Maryland, Baltimore County, United States

Ying Zou , Queen's University, Canada

Zhen Ming Jack Jiang , York University, Canada

ETT Position Paper Program Committee

Ying (Jenny) Zou (PC Chair) , Queen's University, Canada

Ken Wong (PC Chair) , University of Alberta, Canada

PC members for ETT

Bram Adams , MCIS, École Polytechnique de Montréal, Canada

Shihong Huang , Florida Atlantic University, United States

Norha M. Villegas , Icesi University, Cali, Colombia

Hausi A. Muller , University of Victoria, Canada

Chung-Horng Lung , Carleton University, Canada

Kamran Sartipi , University of Ontario Institute of Technology, Canada

Michael Smit , Dalhousie University, Canada

Thomas Dean , Queen's University, Canada

Iman Keivanloo , Queen's University, Canada

Eric Yu , University of Toronto, Canada

Pat Martin , Queen's University, Canada

IBM

CASCON Website Archives

Find below PDF versions of previous CASCON websites

Cascon 2014 (82274KB)

Cascon 2013 (2MB) (2188KB)

Cascon 2012 (567KB) (588KB)

Cascon 2009 (40MB)

Cascon 2008 (20MB)

Cascon 1998-2007 (8MB) (8MB)

CASCON Proceedings Archives

Find below proceedings from previous CASCON events

CASCON 2014
CASCON 2013
CASCON 2012
CASCON 2011
CASCON 2010
CASCON 2009
CASCON 2008
CASCON 2007
CASCON 2006
CASCON 2005
CASCON 2004
CASCON 2003
CASCON 2002
CASCON 2001

Recent News

(NEW!) CASCON 2015 Workshop Materials are now available for Download!

(NEW!) CASCON 2015 Survey is now Available

(NEW!) CASCON 2016, Nov. 8th-10th. Mark you calendar. Program co-chairs for CASCON 2016 are Hausi A. Müller, University of Victoria and Vio Onut, IBM Canada CAS

(NEW!) Recorded keynotes, FoSP sessions and workshop presentations will be available for download after the conference. Stay tuned.

Reminder: CASCON 2015 Opening Keynote Session starts at 9AM on Monday, Nov 2nd, 2015

CASCON 2015 Program Guide is available for download Now

CASCON 2015 Proceedings is available for download Now

Need help on CASCON website? Contact cwillis@ca.ibm.com or casresearch@ca.ibm.com

Follow us through these social outlets:

Facebook

Google+

LinkedIn

Twitter

Related Links

The Four Languages of Innovation: A Model

of Industrial Research and Development (1MB)

Research Partner

Ontario Centres of Excellence (OCE)

Consortium for Software Engineering Research (CSER)

National Research Council Canada (NRC)

Interested in Becoming a Faculty Fellow?

Click Here to Read about the Requirements

.

CASCON Hotel Information

Hotel Name:	Hilton Suites Toronto/Markham Conference Centre & Spa
Hotel Address:	8500 Warden Avenue, Markham, Canada L6G 1A5
Phone Number:	905-470-8500

For directions to the hotel, please click here.

To reserve a hotel room during the conference, please click here.

CASCON 2015 Parking Voucher is Now Available!

Recent News

(NEW!) CASCON 2015 Workshop Materials are now available for Download!

(NEW!) CASCON 2015 Survey is now Available

(NEW!) CASCON 2016, Nov. 8th-10th. Mark you calendar. Program co-chairs for CASCON 2016 are Hausi A. M[illegible]ller, University of Victoria and Vio Onut, IBM Canada CAS

(NEW!) Recorded keynotes, FoSP sessions and workshop presentations will be available for download after the conference. Stay tuned.

Reminder: CASCON 2015 Opening Keynote Session starts at 9AM on Monday, Nov 2nd, 2015

CASCON 2015 Program Guide is available for download Now

CASCON 2015 Proceedings is available for download Now

Need help on CASCON website? Contact cwillis@ca.ibm.com or casresearch@ca.ibm.com

Follow us through these social outlets:

Facebook

Google+

LinkedIn

Twitter

Related Links

The Four Languages of Innovation: A Model of Industrial Research and Development

Research Partner

Ontario Centres of Excellence (OCE)

Consortium for Software Engineering Research (CSER)

National Research Council Canada (NRC)

Interested in Becoming a Faculty Fellow?

Click Here to Read about the Requirements

.