CASCON 2018

Think: Cognitive Computing, Big Data, Cloud, Security and Privacy

October 29 - 31

CASCON News

Download the CASCON 2018 Program Guide

See the CASCON 2018 Proceeding in ACM

Get certain Workshops materials

Call for Papers for CASCON 2019

Give Feedback

Share the video:

28th Annual International Conference on Computer Science and Software Engineering

CASCON is a premier industrial and academic conference for advanced studies in computer science and software engineering sponsored by the IBM Canada Laboratory. CASCON attracts software developers, researchers, innovators, technologists and decision makers from academia, industry, and government who come together to learn about technology trends, present papers, participate in workshops, and exhibit prototypes and solutions. The CASCON program features thought-provoking keynotes, original research papers, stimulating workshops, and an engaging technology expo. Conference proceedings are published in the ACM Digital Library.

Give Feedback

2018 KEYNOTE SPEAKERS

Data Science for Everyone

The Quantum Computing Landscape

Encryption, Key Management,

It Only Took 60 Years to Solve Artificial Intelligence

Monday, October 29th
14:00 - 15:00

Steven Astorino

VP of Development, Hybrid Cloud,
z Analytics and Canada Lab Director at IBM

Click here to view the keynote.

Tuesday, October 30th
11:00 - 12:00

Jessica Pointing

Quantum Computing Researcher at Stanford, Harvard, MIT

Click here to view the keynote.

and Quantum Computing

Tuesday, October 30th
14:00 - 15:00

Walid Rjaibi

CTO, Data Security at IBM

Click here to view the keynote.

Wednesday, October 31st
11:00 - 12:00

John Tsotsos

Distinguished Research Professor of Vision Science at York University's Lassonde School of Engineering, and Director of the Centre for Innovation in Computing at Lassonde

Click here to view the keynote.

Read More

Give Feedback

SHOWCASE YOUR RESEARCH. SHARE YOUR KNOWLEDGE.

We cordially invite research and position papers as well as workshop and expo proposals in the broad areas of computer science, software engineering, software applications and data science in general, including but not limited to cognitive computing, big data analytics, cloud computing, security and privacy, Internet of Things (IoT), software-as-a-service, adaptive systems, service systems, database technologies, compiler technologies, machine learning, mobile application frameworks and platforms, blockchain, DevOps, and business intelligence.

We also encourage papers, workshop, and expo proposals that describe applications including but not limited to the following areas: healthcare, finance, smart cities, supply chain, smart commerce, smart grid and energy, connected and driverless cars, smart design, wearable computing, and social computing.

We invite submissions that that focus on innovation and innovative use of products, platforms and solutions including but are not limited to Bluemix, Cloud Foundry, Analytics, Watson, Cognos, PureApplication, IRIS, Weather, IoT Platform, RFID, DB2, Rational, Websphere, Worklight, Jazz, Lotus.

Give Feedback

Technical Papers

Submit your paper and showcase your research findings to a variety of attendees from different backgrounds and specialties. If your paper is accepted, it will be showcased in the conference proceedings published by CASCON and included in the ACM Digital Library. Visit the Papers section to find out more about the submission steps and requirements.

Workshops

You will have the opportunity to present, discuss, and debate ideas, challenges, emerging technologies and project directions in an interdisciplinary environment.

Expo

Best Paper Awards

Our famous technology showcase is interactive, innovative, and includes diverse exhibits of research interests, technologies, joint projects, and collaboration opportunities.

CASCON recognizes the best technical contributions of the event in terms of originality, clarity, and potential impact with two awards: Best Paper and Best Student Paper. To be eligible for the Best Student Paper award, a student must have primarily authored the paper, and student(s) at the time must have done the work described.

Give Feedback

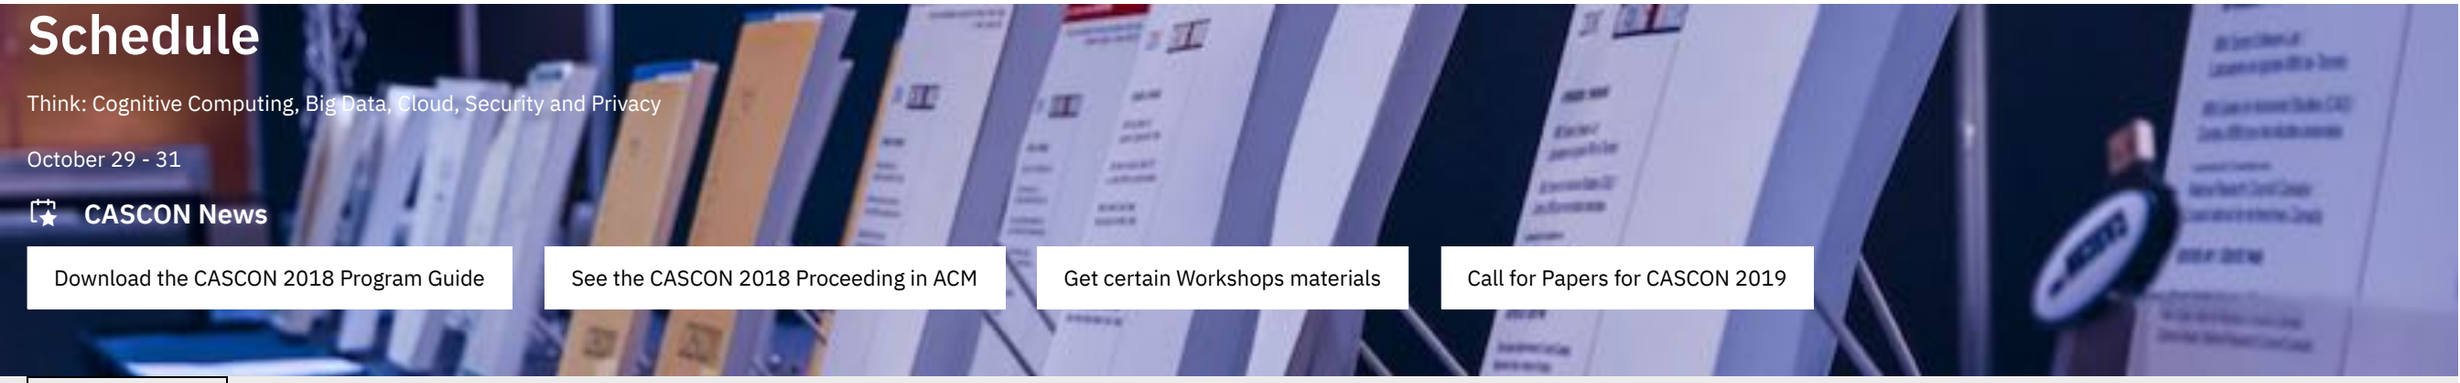

Schedule

Think: Cognitive Computing, Big Data, Cloud, Security and Privacy

October 29 - 31

CASCON News

Download the CASCON 2018 Program Guide | See the CASCON 2018 Proceeding in ACM | Get certain Workshops materials | Call for Papers for CASCON 2019

Full Schedule | Monday, Oct 29 | Tuesday, Oct 30 | Wednesday, Oct 31

	Monday, October 29		Tuesday, October 30		Wednesday, October 31	
8:00	8:00 - 17:00 Registration 8:00 - 11:00 Expo Setup	8:00 - 9:00 Breakfast	8:00 - 17:00 Registration 8:00 - 11:00 Expo Setup	8:00 - 9:00 Breakfast	8:00 - 8:30 Registration 8:00 - 8:30 Expo Setup	8:00 - 9:00 Breakfast
8:15						
8:30	8:30 - 10:00 Regular Paper Track		8:30 - 10:00 Regular Paper Track	8:30 - 14:30 IBM STEM 4 Girls Workshop	8:30 - 10:45 Expo	8:30 - 10:45 Workshops
8:45						
9:00						
9:15						
9:30						
9:45						
10:00	Break		Break			
10:15	10:15 - 10:45 Position Papers		10:15 - 10:45 Position Papers			
10:30						
10:45	Nutrition Break		Nutrition Break		Nutrition Break	
11:00						

11:15	11:00 - 13:00 **Expo**	11:00 - 12:00 Keynote #2 **Jessica Pointing** Harvard/MIT	11:00 - 12:00 Keynote #4 **John K. Tsotsos** York University
11:30			
11:45			
12:00	12:00 - 13:00 Lunch	12:00 - 13:30 **Expo**	12:00 - 13:30 **Expo**
12:15		12:00 - 13:00 Lunch	12:00 - 13:00 Lunch
12:30			
12:45			
13:00	13:00 - 13:30 Opening Ceremony		
13:15			
13:30	13:30 - 13:45 CAS Picture	13:30 - 13:40 Paper Awards	13:30 - 13:45 Expo Awards 13:45 - 14:00 CASCON 2019 Announcement
13:45	13:45 - 14:00 CAS Awards	13:40 - 14:00 **Most Influential Paper (MIP)**	
14:00	14:00 - 15:00 Keynote #1 **Steven Astorino** IBM	14:00 - 15:00 Keynote #3 **Walid Rjaibi** IBM	Break
14:15			14:15 - 15:30 **Workshops**
14:30			
14:45			
15:00	Nutrition Break	Nutrition Break	
15:15	15:15 - 17:15 **Workshops**	15:15 - 17:15 **Workshops**	
15:30			Nutrition Break
15:45			15:45 - 17:30 **Workshops**
16:00			
16:15			
16:30			
16:45			
17:00			
17:15	Break	Break	
17:30	17:30 - 19:00 **Expo Reception** **Posters / Short Talks**	17:30 - 19:00 **Expo Reception** **Posters / Short Talks**	CASCON Ends at 17:30
17:45			
18:00			
18:15			

Give Feedback

18:30		
18:45		
19:00	Expo Closes at 19:00	Expo Closes at 19:00

Give Feedback

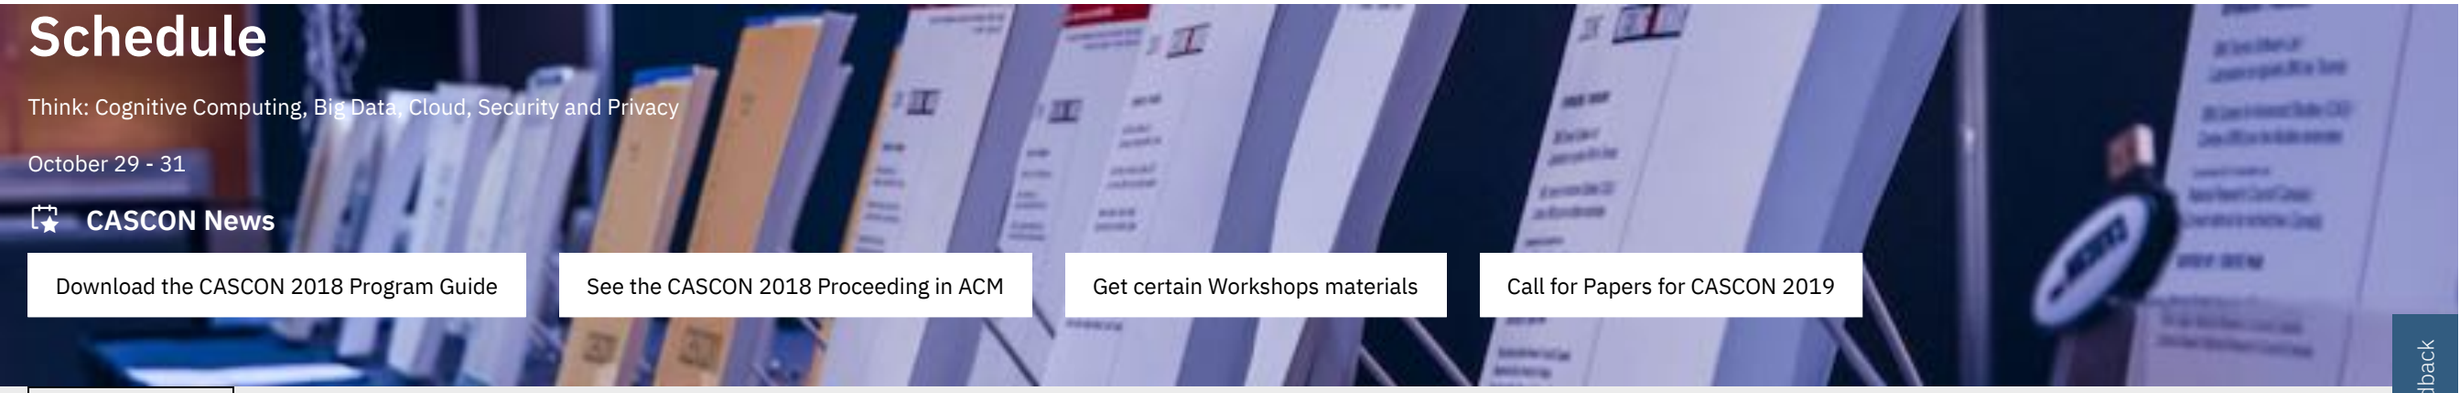

Schedule

Think: Cognitive Computing, Big Data, Cloud, Security and Privacy

October 29 - 31

CASCON News

Download the CASCON 2018 Program Guide | See the CASCON 2018 Proceeding in ACM | Get certain Workshops materials | Call for Papers for CASCON 2019

Full Schedule | **Monday, Oct 29** | Tuesday, Oct 30 | Wednesday, Oct 31

	Monday, October 29	
8:00	8:00 - 17:00 Registration 8:00 - 11:00 Expo Setup	8:00 - 9:00 Breakfast
8:15		

8:30	8:30 - 10:00
	Regular Paper Track
	• Evergreen : Model-Driven Engineering
	  • The Impact of Design and UML Modeling on Codebase Quality and Sustainability
	  • UML-Driven Automated Software Deployment
	  • Automating the Detection of Third-Party Java Library Migration At The Function Level
	• Holly-Buttermut : Software Product Lines
	  • Assuring the runtime behavior of self-adaptive cyber-physical systems using feature modeling
	  • Reducing Variability of Technically Related Software Systems in Large-Scale IT Landscapes
	• Jasmine : Big Data Analytics	

8:45		
9:00		
9:15		
9:30		
9:45		

- A Case Study of Spark Resource Configuration and Management for Image Processing Applications
- Detecting Communities in Social Networks Using Concept Interestingness
- Feature engineering in Big Data for detection of information systems misuse

- Primrose : Cognitive Computing
  - A Competitive Platform for Continuous Programming Skill Enhancement
  - Graphics Programming in Elm Develops Math Knowledge & Social Cohesion
  - Natural Language Question Answering in the Financial Domain

10:00 Break

10:15 - 10:45
Position Papers

10:15
- Evergreen : Microservices
  - Adaptation as a Service
  - Uncertainty Quantification-as-a-Service
- Jasmine : Adaptive Systems

10:30
  - A DevOps Framework for Quality-Driven Self-Protection in Web Software Systems
  - Ontology Driven Temporal Event Annotator mHealth Application Framework
- Primrose : Compiler Development
  - A Survey of Ahead-of-Time Technologies in Dynamic Language Environments
  - All Timescale Window Co-occurrence

10:45 Nutrition Break

11:00 11:00 - 13:00
Expo

11:15
- B1: Document Signing on Blockchain

11:30
- B2: PowerAI - Democratizing AI for All Research Communities
- B3: Data Science 1 - Competency Model

11:45
- B4: Performance and Security Enhanced Containerized Analytics for IoT Systems

12:00
- B5: Randomized Offloading Algorithms for Mobile Cloud Computing
- B6: Exploring Challenges In Adoption of Continuous Delivery

12:15
- B7: Container Technologies! Which One Best Suits Eclipse IoT Deployments?

12:30
- B8: Fall as an irregular pattern in IoT streaming data

12:45
- B9: DevOps Round-trip Engineering: Traceability from Dev to Ops and Back
- B10: Installing features to the Open Liberty kernel with Maven and Gradle
- B11: A Hierarchical Rule-based Security Management System for Data-Intensive Applications
- B12: Continuous Delivery with Open Toolchains
- B13: Ahead-Of-Time Compilation of WebAssembly Using Eclipse OMR

12:00 - 13:00
Lunch

Give Feedback

	• B14: Requirements-based Framework for Evaluating Trust in Service-oriented Systems • B15: Tracking Unhandled Rejections in Node.js • B16: A semi-automated Framework for Migrating Web applications from SQL to NoSQL database • B17: Feature engineering in Big Data for the detection of information systems misuse • B18: Composing Web Services Using a Multi-Agent Framework • B19: Detection of Feature Interaction in Dynamic Scripting Languages • B20: Enabling a highly-scalable, hybrid cloud-based microservices architecture • B21: The Power of Accelerators for Cognitive Workloads • B22: Predicting Cross-Sell Opportunities in Financial Services Using Deep Learning • B23: Attention models for gene promoter sequence characterization • B24: Using AP-TED to Detect Phishing Attack Variations
13:00	13:00 - 13:30 Opening Ceremony
13:15	
13:30	13:30 - 13:45 CAS Picture
13:45	13:45 - 14:00 CAS Awards
14:00	14:00 - 15:00 Keynote #1 **Steven Astorino** IBM
14:15	
14:30	
14:45	
15:00	Nutrition Break
15:15	15:15 - 17:15 - **Workshops** • Building AI models using IBM Watson Studio *(HandsOn, in Conf Center 2)* • Build a cognitive serverless Slack app with IBM Cloud Functions & IBM Watson API *(HandsOn, in Holly-Butternut)* • Build better APIs with the next generation of API testing and monitoring *(HandsOn, in Orchid)* • Come learn how to deploy Open Liberty applications using Docker, Kubernetes, Helm and MicroProfile! *(HandsOn, in Violet)* • Best Practices and Lessons Learned in Microservices *(Panel, in Evergreen)* • Data-driven medicine: promise and challenges *(Panel, in Primrose)* • 13th Workshop on Challenges For Parallel Computing *(Speakers, in Conf Center 1)* • The 3rd International Workshop on Dew Computing *(Speakers, in Elm2)* • IBM Academic Skills Academy - Syllabi and Other Things *(Speakers, in Jasmine)*
15:30	
15:45	
16:00	
16:15	
16:30	
16:45	
17:00	
17:15	Break
17:30	17:30 - 19:00 **Expo Reception**

Give Feedback

17:45	**Posters / Short Talks**
18:00	Short Talks start from 17:30 (5 mins each). • B1: Document Signing on Blockchain
18:15	• B15: Tracking Unhandled Rejections in Node.js
18:30	• B16: A semi-automated Framework for Migrating Web applications from SQL to NoSQL database
18:45	• B4: Performance and Security Enhanced Containerized Analytics for IoT Systems • B5: Randomized Offloading Algorithms for Mobile Cloud Computing • B6: Exploring Challenges In Adoption of Continuous Delivery • B7: Container Technologies! Which One Best Suits Eclipse IoT Deployments? • B8: Fall as an irregular pattern in IoT streaming data Short Talks start from 18:10 (10 mins each). • S1: Db2 on Cloud Free Forever Cloud Database • S2: User Research: Mapping the end-to-end experience of a System Administrator's deployment • S3: Assessing Databases for GDPR Compliance • S4: A Dynamic Knowledge Management System for Decision Support • S5: Distribution and Severity of Test Smells: An exploratory study
19:00	Expo Closes at 19:00

Give Feedback

Schedule

Think: Cognitive Computing, Big Data, Cloud, Security and Privacy

October 29 - 31

CASCON News

Download the CASCON 2018 Program Guide

See the CASCON 2018 Proceeding in ACM

Get certain Workshops materials

Call for Papers for CASCON 2019

Give Feedback

Full Schedule | Monday, Oct 29 | **Tuesday, Oct 30** | Wednesday, Oct 31

	Tuesday, October 30	
8:00	8:00 - 17:00 Registration 8:00 - 11:00 Expo Setup	8:00 - 9:00 Breakfast
8:15		

8:30	8:30 - 10:00
	Regular Paper Track
	• Evergreen : Resource Mgmt.
	  • Feasibility of Internal Object Pools for Reduced Memory Management
	  • Hardware/Software CoDesign for Mathematical Function Acceleration
	  • Persistent Memory Storage of Cold Regions in the OpenJ9 Java Virtual Machine
	• Holly-Butternut : Security and Privacy
	  • Empirical Vulnerability Analysis of Automated Smart Contracts Security Testing on Blockchains
	  • Just-in-time Detection of Protection-Impacting Changes on WordPress and MediaWiki
	  • Scalable Practical Byzantine Fault Tolerance with Short-Lived Signature Schemes
	• Jasmine : Cloud Systems Mgmt.
	  • Design and Implementation of Loss Mitigation in Spot Instances
	  • Evaluating Efficiency, Effectiveness and Satisfaction of AWS and Azure from the Perspective of Cloud Beginners
	  • Node.js Scalability Investigation in the Cloud
		8:30 - 14:30
		IBM STEM
		4 Girls
		Workshop

8:45		
9:00		
9:15		
9:30		
9:45		

Give Feedback

	• Primrose : Machine Learning
	  • A Context-Aware Machine Learning-based Approach
	  • Evaluating Music Mastering Quality Using Machine Learning
	  • (Semi)Automatic Construction of Access-Controlled Web Data Services	

10:00	Break	

10:15	10:15 - 10:45
	Position Papers
	• Jasmine : Blockchain
	  • Challenges and solutions on architecting Blockchain Systems
	  • Powering Software Sustainability with Blockchain	

10:30	• Primrose : Healthcare
	  • Ischemic Brain Stroke Detection using EEG signals
	  • Predictive Analytics in Healthcare: Epileptic Seizure Recognition	

10:45	Nutrition Break	
11:00	11:00 - 12:00 Keynote Speaker #2 **Jessica Pointing** Harvard/MIT	
11:15		
11:30		
11:45		
12:00	12:00 - 13:30 **Expo**	
12:15	• B1: Foodie: Conversational AI for the smart kitchen	

Give Feedback

12:30 12:45	• B2: Privacy-aware Efficient Visual Recognition Services for Smart Laboratories • B3: Data Science 2 - Bootcamp Materials • B4: In-Depth Evaluation of Task Parallelism in Node.js • B5: Supporting Software Variability in Eclipse OMR • B6: Fear and Disgust during the 2015 Canadian Election: Automated Analysis with Watson • B7: ROAD - Realtime Outlier and Anomaly Detection for IBM Cloud • B8: Reconstruct the genealogy of phishing kits source code modifications • B9: Machine Learning to Provide Intelligent Digital Health for Underserved Regions • B10: Detecting Misuses of Crypto APIs • B11: Pressure Sensitive Mats for Patient Monitoring in the NICU • B12: Susereum - A Blockchain Platform for Sustainable Software Systems • B13: Localizing Error Prone Code Using Software Analytics • B14: Eclipse OMR on the Embedded AArch64 Platform • B15: (Semi)Automatic Construction of Access-Controlled Web Data Services • B16: Design Patterns for Machine Learning Solutions • B17: Open Liberty and Eclipse MicroProfile: Foundation for Microservices • B18: Node.js Scalability Investigation in the Cloud • B19: Feasibility of Internal Object Pools to Reduce Memory Management Activity • B20: Cold Object Identification, Sequestrating and Revitalization	12:00 - 13:00 Lunch

Give Feedback

	• B21: Persistent Memory Storage of Cold Regions in the OpenJ9 Java Virtual Machine • B22: A novel approach to generational garbase collection • B23: Math function performance acceleration via integrated HW and algorithm design • B24: Automating the Detection of Third-Party Java Library Migration At The Function Level	
13:00		
13:15		
13:30	13:30 - 13:40 Paper Awards	
13:45	13:40 - 14:00 **Most Influential Paper (MIP)**	
14:00	14:00 - 15:00 Keynote #3 **Walid Rjaibi** IBM	
14:15		
14:30		
14:45		
15:00	Nutrition Break	
15:15	15:15 - 17:00 - **Workshops** • Blockchain Fundamentals and Development Platforms *(HandsOn, in Conf Center 2)* • Modernize digital applications with Microservices management using the Istio service mesh *(HandsOn, in Holly-Butternut)* • Hands-On: Easy Microservices Application Development with Microclimate *(HandsOn, in Orchid)* • IBM Voice Agent with Watson *(HandsOn, in Violet)* • Third Annual Workshop on Data-Driven Knowledge Mobilization *(Panel, in Jasmine)*	
15:30		
15:45		
16:00		
16:15		

Give Feedback

16:30	• The 10th CASCON Workshop on Cloud Computing *(Speakers, in Conf Center 1)* • CASCON Workshop on Developing Big Data Applications and Services - BDAS 2018 *(Speakers, in Elm2)*
16:45	• The Best of IBM Innovation: Advancements through Overcoming Technological Uncertainties *(Speakers, in Evergreen)*
17:00	• Large-Scale Multilevel Streaming Data Analytics *(Speakers, in Primrose)*
17:15	Break
17:30	17:30 - 19:00 **Expo Reception** **Posters / Short Talks** Short Talks start from 17:30 (5 mins each). • B1: Foodie: Conversational AI for the smart kitchen • B2: Privacy-aware Efficient Visual Recognition Services for Smart Laboratories • B16: Design Patterns for Machine Learning Solutions • B4: In-Depth Evaluation of Task Parallelism in Node.js • B5: Supporting Software Variability in Eclipse OMR • B6: Fear and Disgust during the 2015 Canadian Election: Automated Analysis with Watson • B7: ROAD - Realtime Outlier and Anomaly Detection for IBM Cloud • B8: Reconstruct the genealogy of phishing kits source code modifications • B9: Machine Learning to Provide Intelligent Digital Health for Underserved Regions • B10: Detecting Misuses of Crypto APIs Short Talks start from 18:20 (10 mins each). • S1: Discovering business rules automatically from legacy applications • S2: REST APIs with LoopBack 4 and OpenAPI 3 • S3: Dynamic Large-Scale Graph Processing over Data Streams
17:45	
18:00	
18:15	
18:30	
18:45	

Give Feedback

	• S4: A Multilevel Streaming Data Analytics Infrastructure for Predictive Analytics
19:00	Expo Closes at 19:00

Give Feedback

Schedule

Think: Cognitive Computing, Big Data, Cloud, Security and Privacy

October 29 - 31

CASCON News

Download the CASCON 2018 Program Guide

See the CASCON 2018 Proceeding in ACM

Get certain Workshops materials

Call for Papers for CASCON 2019

Give Feedback

Full Schedule	Monday, Oct 29	Tuesday, Oct 30	**Wednesday, Oct 31**

	Wednesday, October 31	
8:00	8:00 - 8:30 Registration 8:00 - 8:30 Expo Setup	8:00 - 9.00 Breakfast
8:15		
8:30	8.30 - 10:45 **Expo** • B1: Model based Interference-aware Management System for Microservice Architectures • B2: Modeling AUTOSAR Implementations in Simulink • B3: Data Science 3 - Student Success Projects • B4: Accessibility in the DevOps Era • B5: Space Sentry Challenge: A game built with IoT devices and Microservices • B6: Comparative Analysis of Cloud Console Systems • B7: REST APIs with LoopBack 4 and OpenAPI 3 • B8: Interactive Visualization to Engage End-users in Prescriptive Analysis • B9: Finding Missing Labels for Large Industrial Datasets • B10: Synthesizing Realistic Data for Machine Learning • B11: Watch Phishing Evolution Over time • B12: Easy Cloud-Native Development with Microclimate • B13: A Machine Learning Approach to Test Case Prioritization • B14: Optimization Driven Inlining • B15: Semantic-Aware Disambiguation for Entity Resolution • B16: Create data-driven visual reports in Reactive Reports, a framework built for big data • B17: OptImatch: Semantic-Web System for Automated Query Tuning • B18: Third Party Container Tests	8:30 - 10:45 **Workshops** • Building Microservices in a Cloud-Native World using Eclipse MicroProfile and Open Liberty *(HandsOn, in Orchid)* • Refine, restructure and make sense of data visually, using IBM Watson Studio *(HandsOn, in Violet)* • 2nd Workshop on DevOps and Software Analytics for Continuous Engineering and Improvement *(Speakers, in Jasmine)* • Practical Machine Learning with Python on DSX *(HandsOn, in Conf Center 2)* • Introduction to the IBM Q experience and Quantum Computing *(HandsOn, in Holly-Butternut)* • Compiler-Driven Performance Workshop *(Speakers, in Conf Center 1)* • Distributed Ledgers and Blockchain: Concepts and Applications *(Speakers, in Evergreen)* • 2nd Workshop on Advances in Open Runtime Technology for Cloud Environments *(Speakers, in Primrose)*
8:45		
9:00		
9:15		
9:30		

Give Feedback

	• B19: Runtime Automatic Parallelization of JVM Applications • B20: Variability Aware analysis using Clang • B21: Approximation Algorithms for Instruction Scheduling and Rapid Prototyping in Coconut • B22: Static Dataflow Analysis for OpenMP Code Sanitization • B23: Plasma Focused Ion Beam Curtaining Artefact Correction by Fourier-Based Linear Model • B24: FLASH: Multithreaded Fault Localization	
9:45		
10:00		
10:15		
10:30		
10:45	Nutrition Break	
11:00	11:00 - 12:00 Keynote #4 **John K. Tsotsos** York University	
11:15		
11:30		
11:45		
12:00	12:00 - 13:30 **Expo Continued**	12:00 - 13:00 Lunch
12:15		
12:30		
12:45		
13:00		

Give Feedback

13:15		
13:30	13:30 - 13:45 Expo Awards 13:45 - 14:00 CASCON 2019 Announcement	
13:45		
14:00	Break	
14:15	14:15 - 17:30 - **Workshops** • IBM Security Guardium Analyzer Bootcamp *(HandsOn, in Orchid)* • Deriving Client Insights in the Financial Sector *(HandsOn, in Violet)* • iCity - Big Data and Visualization Urban Transportation Strategies *(Speakers, in Jasmine)* • Practical Machine Learning with Python on DSX *(HandsOn, in Conf Center 2)* • Introduction to the IBM Q experience and Quantum Computing *(HandsOn, in Holly-Butternut)* • Compiler-Driven Performance Workshop *(Speakers, in Conf Center 1)* • Distributed Ledgers and Blockchain: Concepts and Applications *(Speakers, in Evergreen)* • 2nd Workshop on Advances in Open Runtime Technology for Cloud Environments *(Speakers, in Primrose)*	
14:30		
14:45		
15:00		
15:15		
15:30	Nutrition Break	
15:45	**Workshops Continued**	
16:00		
16:15		
16:30		
16:45		
17:00		

Give Feedback

17:15	
17:30	CASCON Ends at 17:30

Give Feedback

Speakers

Think: Cognitive Computing, Big Data, Cloud, Security and Privacy

October 29 - 31

CASCON News

Download the CASCON 2018 Program Guide

See the CASCON 2018 Proceeding in ACM

Get certain Workshops materials

Call for Papers for CASCON 2019

Give Feedback

Steven Astorino

VP of Development, Hybrid Cloud, z Analytics and Canada Lab Director at IBM

Data Science for Everyone

Monday, October 29th, 14:00 - 15:00

Click **HERE** to see the keynote.

Share the keynote:

Short Abstract: Although an empirical approach to data is as old as the scientific method, the data driven transformation of business and society has initiated a renaissance, with data science and analytics emerging as a key engine of wealth and prosperity for the 21st century. Will better tools and better training satisfy the exploding demand? Should data science take its place with physics, chemistry and biology as a basic science everyone needs to learn? What opportunities will be created, and how will we take advantage of them?

Biography: Steven is a development executive with proven transformational leadership and expertise in leading large enterprise development teams and an experienced change agent with a proven track record to drive improvements and efficiency. Most recently he has been leading and driving the Machine Learning and Data Science strategy for IBM's analytics group and has delivered key

Give Feedback

strategic solutions including IBM Machine Learning and Data Science Experience.

Early in his career, Steven has spent several years working with network testing technologies for the telecom Industry and played a key role in providing VoIP testing solutions and has extensive experience with data technologies such as data movement and transformation from his days at DataMirror, to Database technologies such as DB2 LUW, DB2 for z/OS, IBM's latest Data and Analytics platform enabling data science and machine learning.

He has a Bachelor's Degree in Computer Science with a strong business background.

Jessica Pointing

Quantum Computing Researcher at Stanford, Harvard, MIT

The Quantum Computing Landscape

Tuesday, October 30th, 11:00 - 12:00

Click **HERE** to see the keynote.

Share the keynote:

Give Feedback

Short Abstract: Quantum computers have the potential to advance many other fields, such as machine learning, medicine, and energy systems. We could solve problems on quantum computers that are not possible to solve on our current classical computers in our lifetime. What are the potential applications of quantum computing? How do quantum computers actually work? How do we build quantum computers? How much progress have we made? Jessica Pointing will unpack the topic of quantum computing with a fascinating primer.

Biography: Jessica Pointing has conducted quantum computing research at Stanford, Harvard and MIT. She is currently a PhD student specializing in quantum computing at Stanford and has been awarded the Knight-Hennessy Fellowship. She completed her Bachelor's degree in Physics and Computer Science, with high honors and as a member of the academic honors society Phi Beta Kappa, from Harvard University after spending her first two years of university studies at the Massachusetts Institute of Technology (MIT). She also founded the Harvard College Quantum Computing Association. Jessica was the winner (audience vote) for the Quantum Matters Science Communication Competition and has been invited to talk about quantum computing at conferences, including IBM's conference in Las Vegas, the Chief Digital Officers conference in New York City and the Neo conference in Arizona. Jessica has interned as a software engineer at Google, management consultant at McKinsey & Company, investment banker at Goldman Sachs, and strategist at Morgan Stanley. She is the founder and editor-in-chief of Optimize Guide, a blog about life optimization, and founded the Now Know Organization, which received funding from Google to teach technology to young people. She has been awarded the McKinsey Women's Impact Award, MIT Award for Distinguished Achievement in Leadership, and has been named a Google Anita Borg Scholar and John Harvard Scholar (top 5% of class). She is also a Microsoft Scholar, Palantir Scholar, Adobe Research Scholar, Morgan Stanley Scholar, Goldman Sachs Scholar, Society of Women Engineers Scholar and Society of Geophysicists Scholar.

Give Feedback

Walid Rjaibi

CTO, Data Security at IBM

Encryption, Key Management, and Quantum Computing

Tuesday, October 30th, 14:00 - 15:00

Click **HERE** to see the keynote.

Share the keynote:

Short Abstract: Encryption is a fundamental technology for data protection. While Transport Layer Security (TLS) is widely accepted as the standard solution for encrypting data in transit, no single solution has achieved similar status for encrypting data at rest. We will review the most common strategies for encrypting data at rest and the use cases where each strategy is most appropriate. We will also discuss how exactly quantum computers will affect data encryption and how those concerns are being addressed.

Biography: Walid Rjaibi is Chief Technology Officer (CTO) for Data Security at IBM. As a Technical Executive and Distinguished Engineer, he drives the technical strategy and architecture for data security across products and cloud services (Guardium Database Activity Monitoring, Guardium Data Encryption, Security Key Lifecycle Manager, Guardium Multi-Cloud Data Encryption, and IBM Key

Give Feedback

Protect). Prior to his current role, he held several technical and management roles within IBM including Research Staff Member at the Zurich Research Lab, Security Architect for DB2 LUW, and Chief Security Architect for IBM Analytics. His Data Security work resulted in over 20 granted patents and several publications in leading scientific and academic conferences such as the international conference on Very Large Databases (VLDB), the International Conference on Data Engineering (ICDE), and the international conference on security and cryptography (SECRYPT). Walid is also a frequent speaker at industrial conferences such as IBM World of Watson, Interconnect, Security Congress, and the International DB2 User Group (IDUG).

Specialties: Information Technology, Software Development, Data Security, Application Security, Cloud Security, Security Compliance, Secure Engineering Deployment, Database Technology, Leadership, Staff Management and Team Building, Client Partnering

---

John Tsotsos

Distinguished Research Professor of Vision Science at York University's Lassonde School of Engineering, and Director of the Centre for Innovation in Computing at Lassonde

It Only Took 60 Years to Solve Artificial Intelligence - That Wasn't so Hard, Was it?

Wednesday, October 31st, 11:00 - 12:00

Click **HERE** to see the keynote.

Give Feedback

Share the keynote:

Short Abstract: It has been 62 years since the founding Dartmouth AI workshop. In the 1960's, researchers at MIT believed that developing a computer program to simulate a significant part of human vision could be accomplished by employing several undergraduate students over a summer. They quickly found out they were mistaken. Fast forward to the present day. AI is everywhere, solving all manner of hard problems and triggering debates about its ethics and societal impact. In contrast, scientists studying human (and animal) vision and intelligence have made no corresponding claims for breakthroughs in their understanding of human intelligence. Although a great deal has been learned in the past 60 years to be sure, each major discovery seems to emphasize how much remains unknown. Understanding the brain and how intelligent behaviour is produced remains a major challenge. Noting that the current AI successes are often claimed to be due to mimicking brain processes, is there a disconnect here? This presentation will suggest that there has been some tacit, gradual moving of the goalposts taking place and that there is indeed a disconnect. An understanding of the brain and behavioural sciences reveals many directions for future AI research, and importantly, how far there is still to go. Dramatically, this will require the abandonment of some of the longest held computational standards. Examples from my lab will illustrate.

Biography: John Tsotsos is Distinguished Research Professor of Vision Science in the Dept. of Electrical Engineering and Computer Science and Director of the Centre for Innovation in Computing at Lassonde, at York University. He has Adjunct Professorships in Computer Science and in Ophthalmology and Vision Sciences at the University of Toronto. He received his doctorate in Computer Science from the University of Toronto developing the first computer system to interpret visual motion depicted in digital image sequences, with application to heart motion analysis. After a postdoctoral fellowship in Cardiology at Toronto General Hospital, he joined the University of Toronto on faculty in both Computer Science and in Medicine. In 1980 he founded the highly respected Computer Vision Group at the University of Toronto, which he led for 20 years. He moved to York University in 2000 as Director of the Centre for Vision Research. Under his directorship, the centre was ranked in the top six interdisciplinary vision research organizations in the world. He has been a Canadian Heart Foundation Research Scholar (1981-83), a Fellow of the Canadian Institute for Advanced Research (1985-95) and currently holds the Canada Research Chair in Computational Vision (2003-2024). He has received many awards and honours including several best paper awards, among them a 1987 inaugural Marr Prize citation, the 1997 CITO Innovation Award for Leadership in Product Development, the 2006 Canadian Image Processing and Pattern Recognition Society Award for

Give Feedback

Research Excellence and Service, the 1st President's Research Excellence Award by York University on the occasion of the University's 50th anniversary in 2009, the 2011 Geoffrey J. Burton Memorial Lectureship from the United Kingdom's Applied Vision Association for significant contribution to vision science, has been an ACM Distinguished Speaker, and is an IEEE Fellow. He was elected as Fellow of the Royal Society of Canada in 2010, and was awarded its 2015 Sir John William Dawson Medal for sustained excellence in multidisciplinary research, the first computer scientist to be so honoured. He has co-founded or been a principal of several companies, and holds a number of patents and technology transfer licences. Visiting positions were held at: University of Hamburg, Germany; Polytechnical University of Crete, Greece; Center for Advanced Studies at IBM Canada; INRIA Sophia-Antipolis, France; and, the Massachusetts Institute of Technology, USA. His current research focuses on a comprehensive theory of visual attention in humans. A practical outlet for this theory forms a second focus, embodying elements of the theory into the vision systems of mobile robots.

Give Feedback

Papers

Think: Cognitive Computing, Big Data, Cloud, Security and Privacy

October 29 - 31

CASCON News

Download the CASCON 2018 Program Guide

See the CASCON 2018 Proceeding in ACM

Get certain Workshops materials

Call for Papers for CASCON 2019

Give Feedback

CASCON 2018 Proceeding can be found from **ACM**

Technical papers of CASCON 2018 should be original **research papers** and be at most 12 pages long in ACM format(style ACM_SigConf). CASCON 2018 **experience reports** share experience and lessons gained from the use of technology that can be useful to others and should be at most 12 pages in ACM format. CASCON 2018 **position papers** share vision, highlight technology gaps, or discuss emerging advances and should be at most 8 pages in ACM format. CASCON 2018 will not accept submissions that have been previously published, are in press, or have been submitted elsewhere. Accepted technical papers will be included in the conference proceedings published by CASCON and included in the ACM Digital Library.

Best Paper Awards

Two paper awards - **Best Paper and Best Student Paper** - will 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 student(s) must have primarily authored the paper, must have been a student(s) at the time the work was done, and must have carried out the work described; only the student author(s) receive this award.

Congratulations to the Winners!

The Best Paper Award goes to "Reducing Variability of Technically Related Software Systems in Large-Scale IT Landscapes".

→ Click to read more and share

The Best Student Paper Award is given to "Just-in-time Detection of Protection-Impacting Changes on WordPress and MediaWiki ".

→ Click to read more and share

Give Feedback

Most Influential Paper Award

Since 2010, CASCON has presented a **"Most Influential Paper" (MIP) Award** to a paper published a decade earlier at CASCON. The MIP Award recognizes lasting contributions and impact to theory and practice.

Congratulations to the Winners!

The Most Influential Paper from CASCON 2008 is awarded to "Is it a Bug or an Enhancement? A Text-based Approach to Classify Change Requests".

→ Click to read more and share

Give Feedback

Congratulations to the Authors!

Accepted Papers

This year we received 68 full papers and 23 position papers.
We accepted 23 full papers and 10 position papers.
Acceptance rate for full papers was 33.8% and for position papers was 43.5%.

Thank you for your submissions. See you all at CASCON 2018.

Full Papers | Position Papers

Session Name: Model-Driven Engineering
Day: Monday, October 29th
Time: 8:30 - 10:00
Room: Evergreen
Session Chair: Marios Fokaefs (Polytechnique Montréal)

The Impact of Design and UML Modeling on Codebase Quality and Sustainability

— *Omar Badreddin (University of Texas) ; Khandoker Rahad (University of Texas) ;*

Short Abstract:

The general consensus of researchers and practitioners is that upfront and continuous software design using modeling languages such as UML improve code quality and reliability particularly as the software evolves over time. Software designs and models help in managing the underlying code complexities which are crucial for sustainability. Recently, there has been increasing evidence suggesting broader adoption of modeling languages such as UML. However, our understanding of the impact of using such modeling and design languages remains limited. This paper reports on a study that aims to characterize this impact on code quality and sustainability. We identify a sample of open source software repositories with extensive use of designs and modeling and compare their code qualities with similar code-centric repositories. Our evaluation focuses on various code quality attributes such as code smells and technical debt. We also conduct code evolution analysis over fiveyear period and collect additional data from questionnaires and interviews with active repository contributors. This study finds that repositories with significant use of models and design activities are associated with reduced critical code smells but are also

Give Feedback

associated with increase in non-critical code smells. The study also finds that modeling and design activities are associated with significant reduction in measures of technical debt. Analyzing code evolution over five year period reveals that UML repositories start with significantly lower technical debt density measures but tend to decline over time.

UML-Driven Automated Software Deployment

— *Luis Rivera (Universidad Icesi, University of Victoria) ; Norha Villegas (Universidad Icesi) ; Gabriel Tamura (Universidad Icesi) ; Miguel Jiménez (University of Victoria) ; Hausi Müller (University of Victoria) ;*

Short Abstract:

Software companies face the challenge of ensuring customer satisfaction through the continuous delivery of functionalities and rapid response to quality issues. However, achieving frequent software delivery is not a trivial task. It requires agile and continuous design, development and deployment of existing and new software features. Over time, managing these systems becomes increasingly complex. This complexity stems, in part, from the deployment pipelines and the myriad possible configurations of the software components. Furthermore, software deployment is a timeconsuming and error-prone process, which, even when automated, can lead to configuration errors and cost overruns. In this paper, we address deployment challenges that developers face during continuous delivery and DevOps. Our proposal consists of Urano, a mechanism for automating the deployment process, which uses UML, an interoperable and de facto modeling standard, as a means of specifying a software architecture and its associated deployment. Our approach is based on the model-driven architecture principles to generate executable deployment specifications from user-defined UML deployment diagrams. We extend this kind of diagrams by defining and applying a UML profile that captures the semantics and requirements of the installation, configuration, and update of software components. Thus, enabling more expressive deployment specifications and their automatic realization. To evaluate Urano, we conducted three case studies that demonstrate its potential to effectively automate software deployment processes in industry.

Automating the Detection of Third-Party Java Library Migration At The Function Level

— *Hussein Alrubaye (Rochester Institute of Technology) ; Mohamed Mkaouer (Rochester Institute of Technology) ;*

Short Abstract:

Give Feedback

The process of migrating between different third-party libraries is very complex. Typically, developers need to find functions in the new library that are most adequate in replacing the functions of the retired library. This process is subjective and time-consuming as developers need to fully understand the documentation of both libraries' Application Programming Interfaces, and find the right match between their functions if it exists. In this context, several studies rely on mining existing library migrations to provide developers with by-example approaches for similar scenarios. In this paper, we propose a mining approach that extracts all the manuallyperformed function replacements for a given library migration. Our approach combines the mined function-change patterns with function-related lexical similarity to accurately detect mappings between replacing/replaced functions. Using our enhanced mining process, we perform a comparative study between state-of-art approaches for detecting migration traces at the function level. Our findings have shown its efficiency in accurately detecting migration fragments and it has enhanced the accuracy of state-of-art approaches in finding correct functions changes. We finally provide the community with a dataset of migrations between popular Java libraries, and their corresponding code changes at the function level.

Give Feedback

Session Name: Software Product Lines
Day: Monday, October 29th
Time: 8:30 - 10:00
Room: Holly-Butternut
Session Chair: Alexei Lapouchnian (University of Toronto)

Assuring the runtime behavior of self-adaptive cyber-physical systems using feature modeling

— *Nayreet Islam (University of Ontario Institute of Technology) ; Akramul Azim (University of Ontario Institute of Technology) ;*

Short Abstract:

A self-adaptive cyber-physical system (SACPS) can adjust its behavior and configurations at runtime in response to varying requirements obtained from the system and the environment. With the increasing use of the SACPS in different application domains, such variations are becoming more common. Users today expect the SACPS to guarantee its functional and timing

behavior even in adverse environmental situations. However, uncertainties in the SACPS environment impose challenges on assuring the runtime behavior during system design. Software product line engineering (SPLE) is considered as a useful technique for handling varying requirements. In this paper, we present an approach for assuring the runtime behavior of the SACPS by applying an SPLE technique such as feature modeling. By representing the feature-based model at design time, we characterize the possible adaptation requirements to reusable configurations. The proposed approach aims to model two dynamic variability dimensions: 1) environment variability that describes the conditions under which the SACPS must adapt, and 2) structural variability, that defines the resulting architectural configurations. To validate our approach, the experimental analysis is performed using two case studies: 1) a traffic monitoring SACPS and 2) an automotive SACPS. We demonstrate that the proposed feature-based modeling approach can be used to achieve adaptivity which allows the SACPS to assure functional (defining execution of the correct set of adaptive tasks) and non-functional (defining execution of SACPS in the expected mode) correctness at runtime. The experimental results show that the feature-based SACPS demonstrates significant improvement in terms of self-configuration time, self-adaptation time and scalability with less probability of failure in different environmental situations.

Reducing Variability of Technically Related Software Systems in Large-Scale IT Landscapes

— *Kenny Wehling (Volkswagen AG) ; David Wille (Technische Universität Braunschweig) ; Christoph Seidl (Technische Universität Braunschweig) ; Ina Schaefer (Technische Universität Braunschweig) ;*

Short Abstract:

The number of software systems in a company typically grows with the business requirements. Therefore, IT landscapes in large companies can consist of hundreds or thousands of different software systems. As the evolution of such large-scale landscapes is often uncoordinated, they commonly comprise different groups of related software systems using a common core technology (e.g., Java Web-Application) implemented by a variety of architectural components (e.g., different application servers or databases). This leads to increased costs and higher effort for maintaining and evolving these software systems and the entire IT landscape. To alleviate these problems, the variability of such technically related software systems has to be reduced. For this purpose, experts have to assess and evaluate restructuring potentials in order to take appropriate restructuring decisions. As a manual analysis requires high effort and is not feasible for large-scale IT landscapes, experts face a major challenge. To overcome this challenge, we introduce a novel approach to automatically support experts in taking

Give Feedback

reasonable restructuring decisions. By providing automated methods for assessing, evaluating and simulating restructuring potentials, experts are capable of reducing the variability of related software systems in large-scale IT landscapes. We show suitability of our approach by expert interviews and an industrial case study with architectures of real-world software systems.

Session Name: Big Data Analytics
Day: Monday, October 29th
Time: 8:30 - 10:00
Room: Jasmine
Session Chair: Mark Chignell (University of Toronto)

A Case Study of Spark Resource Configuration and Management for Image Processing Applications

— *Dwight Makaroff (University of Saskatchewan) ; Derek Eager (University of Saskatchewan) ; Winfried Grassmann (University of Saskatchewan) ; Habib Sabiu (University of Saskatchewan) ; Owolabi Adekoya (University of Saskatchewan) ;*

Short Abstract:

The world population is expected to reach an estimated 9.8 billion by 2050, necessitating substantial increases in food production. Achieving such increases will require large-scale application of computer informatics within the agricultural sector. In particular, application of informatics to crop breeding has the potential to greatly enhance our ability to develop new varieties quickly and economically. Achieving this potential, however, will require capabilities for analyzing huge volumes of data acquired from various field-deployed image acquisition technologies. Although numerous frameworks for big data processing have been developed, there are relatively few published case studies that describe user experiences with these frameworks in particular application science domains. In this paper, we describe our efforts to apply Apache Spark to three applications of initial interest within the Plant Pheno- typing and Imaging Research Centre (P2IRC) at the University of Saskatchewan. We find that default Spark parameter settings do not work well for these applications. We carry out extensive performance experiments to investigate the impact of alternative Spark parameter settings, both for applications run individually and in scenarios with multiple concurrently executing applications. We find that optimizing Spark parameter settings is challenging, but can yield substantial performance improvements, particularly with concurrent applications,

Give Feedback

provided that the dataset characteristics are considered. This is a first step towards insights regarding Spark parameter tuning on these classes of applications that may be more generally applicable to broader ranges of applications.

Detecting Communities in Social Networks Using Concept Interestingness

— *Mohamed-Hamza Ibrahim (University of Quebec in Outaouais) ; Rokia Missaoui (UQO - LARIM) ; Abir Messaoudi (UQO - LARIM) ;*

Short Abstract:

One key challenge in Social Network Analysis is to design an efficient and accurate community detection procedure as a means to discover intrinsic structures and extract relevant information. In this paper, we introduce a novel strategy called (COIN), which exploits COncept INterestingness measures to detect communities based on the concept lattice construction of the network. Thus, unlike off-the-shelf community detection algorithms, COIN leverages relevant conceptual characteristics inherited from Formal Concept Analysis to discover substantial local structures. On the first stage of COIN, we extract the formal concepts that capture all the cliques and bridges in the social network. On the second stage, we use the stability index to remove noisy bridges between communities and then percolate (merge) relevant adjacent cliques. Our experiments on several real-world social networks show that COIN can quickly detect communities more accurately than existing prominent algorithms such as Edge betweenness, Fast greedy modularity, and Infomap.

Feature engineering in Big Data for detection of information systems misuse

— *Eduardo Lopez (McMaster University) ; Kamran Sartipi (East Carolina University) ;*

Short Abstract:

The increasing availability of very large volumes of digital data (i.e. Big Data) enables many interesting research streams on a wide variety of phenomena. However, there has been a paucity of Big Data sets in the area of cybersecurity in information systems, as organizations are reluctant to share data that may provide too much unrestricted visibility into their operations. In this study, we explore the use of a real-life, anonymized, very large dataset containing user behavior as captured in log files including both regular usage as well as misuse, typifying the dynamics found in a situation with compromised user credentials. Through the experiment, we validate that the existence of a large user behavior dataset in itself does not necessarily guarantee

Give Feedback

that abnormal behaviors can be found. It is essential that researchers apply deep domain knowledge, critical thinking and practical focus to ensure the data can produce the knowledge required for the ultimate objective of detecting an insider's threat. In this paper we develop, formulate and calculate the features that best represent user behavior in the underlying information systems, maintaining a parsimonious balance between complexity, resource demands and detection effectiveness. We test the use of a classification model that proves the usefulness and aplicability of the features extracted.

Session Name: Cognitive Computing
Day: Monday, October 29th
Time: 8:30 - 10:00
Room: Primrose
Session Chair: Ulrike Stege (University of Victoria)

A Competitive Platform for Continuous Programming Skill Enhancement

— *Jen-Hao Kuo (National Cheng Kung University) ; Tsung-Han Wu (National Cheng Kung University) ; Hong-Bao Ye (National Cheng Kung University) ; Hewijin Christine Jiau (National Cheng Kung University) ;*

Short Abstract:

Enhancing programming skills is the key factor to keep up with current ever-changing technologies in IT industry. Implementing strategies on game-based platform is a common way for programmers to enhance programming skills. However, runtime simulation and game metrics provided by current game-based strategy platforms are ineffective to motivate continuous programming skill enhancement. We propose ELOP, a competitive game-based strategy platform to motivate programmers. ELOP will automatically schedule competitions for programmers, keep competition history, record changes in performance and provide needed personal information for further enhancement. To evaluate the effectiveness of ELOP, we conduct several studies. The result shows that ELOP does motivate programmers in continuous programming practice and enhance their programming skills.

Give Feedback

Graphics Programming in Elm Develops Math Knowledge & Social Cohesion

— *John Zhang (McMaster University) ; Anirudh Verma (McMaster University) ; Chinmay Sheth (McMaster University) ; Christopher Schankula (McMaster University) ; Stephanie Koehl (McMaster University) ; Andrew Kelly (Hamilton-Wentworth District School Board) ; Yumna Irfan (McMaster University) ; Christopher K. Anand (McMaster University) ;*

Short Abstract:

At McMaster University, we have developed a framework for teaching computer science, including curricula and tools (iPad apps: Image 2 Bits and ElmJr; an open-source library GraphicSVG; and a web-based development environment). ElmJr is a projectional editor for Elm, with knowledge of our graphics library. Using ElmJr, children transform programs through contextual menus. As a result, they never see syntax or type errors. Children as young as 10 years old, who have just started learning about syntax in English, can be productive programmers. We will explain how ElmJr is designed to make programming simple for beginners, and how strong typing in Elm and our graphics library is key to making the list of program transformations manageable. We will explain the design and findings of a study of children in 14 classes (grade 4 and 5) in the Hamilton-Wentworth District School Board, who received 12 hours of instruction over eight weeks in ElmJr with the aim of improving their mathematics knowledge. In parallel, another eight grade 6 to 8 classes received at least two hours of instruction in Elm using our web-IDE, culminating in a Wordathon. The Wordathon is designed to reconcile the power of social computing with the need to protect children's privacy. The intermediate classes were challenged to create animations in Elm of assigned words identified by teachers as K-4 core reading vocabulary. Joining the intermediate students in this activity, two high school classes were taught how to create interactive applications in Elm and were challenged to create a reading game using the word animations created by the intermediate children. In all, four games incorporating 408 animations were created, and some of the intermediate students presented the games to primary grades in their schools. Unlike typical network effects, we get a multiplier effect: more word animations make game development more attractive, and more games make animations more attractive. Having a critical mass then attracts the attention of other educators, including in this case the team implementing the board-wide reading strategy. All of this can be accomplished without any identifying information leaving the classroom.

Natural Language Question Answering in the Financial Domain

Give Feedback

— *John Boyer (IBM Canada Ltd.) ;*

Short Abstract:

This paper describes a natural language question answering system focused on answering financial domain questions using a daily updated corpus of financial reports. Financial entity types of interest included company stocks, country bonds, currencies, industries, commodities, and diversified assets. Financial questions of interest included explanatory and factual questions about entities as well as financial outlook for entities. An important architectural divergence emerged between the approach required for answering financial outlook questions versus the approach for answering other financial information questions. The financial domain focus also introduced additional challenges to open domain natural language processing that were addressed in the areas of document ingestion, question classification accuracy, question analysis techniques, speed of machine learning, answer ranking by linguistic confidence versus temporality, and system accuracy assessment.

Session Name: Resource Mgmt.
Day: Tuesday, October 30th
Time: 8:30 - 10:00
Room: Evergreen
Session Chair: Gerhard Dueck (University of New Brunswick)

Feasibility of Internal Object Pools for Reduced Memory Management

— *Konstantin Nasartschuk (University of New Brunswick) ; Kenneth Kent (University of New Brunswick) ; Stephen MacKay (University of New Brunswick) ; Aleksander Micic (IBM Canada Ltd.) ;*

Short Abstract:

Object pools is a widely used software engineering pattern used to reuse object instances without the need of repeated allocation and instantiation. While the benefits of using object pool structures are still present when used in a garbage collected environment, it adds a memory management component to the development process. The paper investigates the

Give Feedback

feasibility of introducing automatically created and maintained object pools for predefined classes. Automatic object pools are implemented and discussed using the GenCon GC and Balanced GC policies.

Hardware/Software CoDesign for Mathematical Function Acceleration

— *Christopher Anand (McMaster University) ; Lucas Dutton (McMaster University) ; Adele Olejarz (McMaster University) ; Robert Enenkel (IBM Canada Ltd.) ; Wolfram Kahl (McMaster University) ;*

Short Abstract:

Many important workloads depend on the efficient computation of elementary functions like square root and logarithm. Accurate computation of these functions is time-consuming, and hard for compilers to schedule, because of conditional execution. These problems are exacerbated by SIMD computation, which does not mix well with conditional execution. Previously, we have outlined how performance can be improved by encapsulating the conditional execution in new instructions. In this paper, we refine this approach to take into account testability, the ability for code to be pipelined, and exploitation of processors with a gather-load instruction. In particular, we look at the decomposition of the previously described instruction pairs into three instructions. The instructions can incorporate table lookups, or complement existing load instructions. The variant which complements existing load instructions is expected to perform as well as the other variants, and is easier to test and to pipeline. This paper presents gate-level details for the instructions required to calculate various logarithm functions, including the circuit depth, count and approximate width. In addition, we highlight the relative complexity of verifying these instructions relative to other known instructions, and outline our strategy for light-weight verification. Finally, we show that this strategy would be expected to produce a doubling of performance on a wide class of processors, using an IBM POWER processor as an example.

Persistent Memory Storage of Cold Regions in the OpenJ9 Java Virtual Machine

— *Scott Young (University of New Brunswick) ; Michael Flawn (University of New Brunswick) ; Kenneth Kent (University of New Brunswick) ; Gerhard Dueck (University of New Brunswick) ; Charlie Gracie (IBM Canada Ltd.) ;*

Short Abstract:

Give Feedback

In this paper an optimization technique for object-oriented language runtimes with automatic memory management is investigated. The technique involves segregating objects into different memory areas, backed by different memory devices, on a per object basis. This technique is compared to operating system paging mechanisms with swap partitions. Two different schemes for determining which objects should be segregated into slower memory are tested and their results are discussed. It has been observed that each technique can be the most or least optimal choice depending on the application.

Session Name: Security and Privacy
Day: Tuesday, October 30th
Time: 8:30 - 10:00
Room: Holly-Butternut
Session Chair: Jeremy Bradbury (University of Ontario Institute of Technology)

Empirical Vulnerability Analysis of Automated Smart Contracts Security Testing on Blockchains

— *Reza Parizi (Kennesaw State University) ; Ali Dehghantanha (University of Guelph) ; Kim Kwang Raymond Choo (University of Texas at San Antonio) ; Amritraj Singh (Kennesaw State University) ;*

Short Abstract:

The emerging blockchain technology supports decentralized computing paradigm shift and is a rapidly approaching phenomenon. While blockchain is thought primarily as the basis of Bitcoin, its application has grown far beyond cryptocurrencies due to the introduction of smart contracts. Smart contracts are self-enforcing pieces of software, which reside and run over a hosting blockchain. Using blockchain-based smart contracts for secure and transparent management to govern interactions (authentication, connection, and transaction) in Internet-enabled environments, mostly IoT, is a niche area of research and practice. However, writing trustworthy and safe smart contracts can be tremendously challenging because of the complicated semantics of underlying domain-specific languages and its testability. There have been high-profile incidents that indicate blockchain smart contracts could contain various code-security vulnerabilities, instigating financial harms. When it involves security of smart contracts, developers embracing the ability to write the contracts should be capable of testing their code, for diagnosing security vulnerabilities, before deploying them to the immutable environments on blockchains.

Give Feedback

However, there are only a handful of security testing tools for smart contracts. This implies that the existing research on automatic smart contracts security testing is not adequate and remains in a very stage of infancy. With a specific goal to more readily realize the application of blockchain smart contracts in security and privacy, we should first understand their vulnerabilities before widespread implementation. Accordingly, the goal of this paper is to carry out a far-reaching experimental assessment of current static smart contracts security testing tools, for the most widely used blockchain, the Ethereum and its domain-specific programming language, Solidity, to provide the first body of knowledge for creating more secure blockchain-based software.

Just-in-time Detection of Protection-Impacting Changes on WordPress and MediaWiki

— *Amine Barrak (Polytechnique Montréal) ; Marc-André Laverdière (Polytechnique Montréal) ; Foutse Khomh (Polytechnique Montréal) ; Le An (Polytechnique Montréal) ; Ettore Merlo (Polytechnique Montréal) ;*

Short Abstract:

Access control mechanisms based on roles and privileges restrict the access of users to security sensitive resources in a multi-user software system. Unintentional privilege protection changes may occur during the evolution of a system, which may introduce security vulnerabilities; threatening user's confidential data, and causing other severe problems. In this paper, we use the Pattern Traversal Flow Analysis technique to identify definite protection differences in WordPress and MediaWiki systems. We analyse the evolution of privilege protections across 211 and 193 releases from respectively WordPress and Mediawiki, and observe that around 60% of commits affect privileges protections in both projects. We refer to these commits as protection-impacting change (PIC) commits. To help developers identify PIC commits just-in-time, we extract a series of metrics from commit logs and source code, and build statistical models. The evaluation of these models revealed that they can achieve a precision up to 73.8% and a recall up to 98.8% in WordPress and for MediaWiki, a precision up to 77.2% and recall up to 97.8%. Among the metrics examined, commit churn, bug fixing, author experiences and code complexity between two releases are the most important predictors in the models. We performed a qualitative analysis of false positives and false negatives and observe that PIC commits detectors should ignore documentation-only commits and process code changes without the comments.

Give Feedback

Scalable Practical Byzantine Fault Tolerance with Short-Lived Signature Schemes

— *Xinxin Fan (IoTeX) ;*

Short Abstract:

The Practical Byzantine Fault Tolerance (PBFT) algorithm is a popular solution for establishing consensus in blockchain systems. The execution time of the PBFT consensus algorithm has an important effect on the blockchain throughput. Digital signatures are extensively used in PBFT to ensure the authenticity of messages during the different phases. Due to the round-based and broadcast natures of PBFT, nodes need to verify multiple signatures received from their peers, which incurs significant computational overhead and slows down the consensus process. To address this issue, we propose an efficient short-lived signature based PBFT variant, which utilizes short-length cryptographic keys to sign/verify messages in PBFT for a short period of time and blockchain-aided key distribution mechanisms to update those keys periodically. We also present efficient algorithms for accelerating the software implementation of the BLS threshold signature scheme. Our extensive experiments with three elliptic curves and two signature schemes demonstrate the efficacy of using short-lived signature schemes for improving the scalability of PBFT significantly.

Give Feedback

Session Name: Cloud Systems Mgmt.
Day: Tuesday, October 30th
Time: 8:30 - 10:00
Room: Jasmine
Session Chair: Farhana Zulkernine (Queen's University)

Design and Implementation of Loss Mitigation in Spot Instances

— *Tasnim Kabir (Bangladesh University of Engineering and Technology) ; A. B. M. Alim Al Islam (Bangladesh University of Engineering and Technology) ;*

Short Abstract:

Spot instances (as provided in Amazon Elastic Compute Cloud, EC2) offer resources at a reduced cost, however, often provide less reliability. This happens as the resources assigned to spot instances can be withdrawn abruptly due to real-time variations in demand and price. A specialized mechanism to deal with this offered reduced cost while maintaining high reliability is desirable to users, as users generally try to minimize their loss maintaining a minimal cost. Such a mechanism is little focused in the literature till now. Therefore, in this paper, we propose a mechanism to mitigate the loss of these spot instances in road to providing high reliability without increasing the computational cost. To do so, first, we apply checkpointing at different stages of the computation. Further, we propose several algorithms to place checkpoints at proper points during a computation. Our proposed algorithms can mitigate the loss by up to 99.9% on an average. We confirm these findings by experimenting over real Amazon spot instance history.

Evaluating Efficiency, Effectiveness and Satisfaction of AWS and Azure from the Perspective of Cloud Beginners

— *Gabriel Costa Silva (Universidade Tecnológica Federal do Paraná) ; Reginaldo Ré (Universidade Tecnológica Federal do Paraná) ; Marco Aurélio Graciotto Silva (Universidade Tecnológica Federal do Paraná) ;*

Short Abstract:

Quality has long been regarded as an important driver of cloud adoption. In particular, quality in use (QiU) of cloud platforms may drive cloud beginners to the cloud platform that offers the best cloud experience. Cloud beginners are critical to the cloud market because they currently represent nearly a third of cloud users. We carried out three experiments to measure the QiU (dependent variable) of public cloud platforms (independent variable) regarding efficiency, effectiveness and satisfaction. AWS EC2 and Azure Virtual Machines are the two cloud services used as representative proxies to evaluate cloud platforms (treatments). Eleven undergraduate students with limited cloud knowledge (participants) manually created 152 VMs (task) using the web interface of cloud platforms (instrument) following seven different configurations (trials) for each cloud platform. Whereas AWS performed significantly better than Azure for efficiency (p-value not exceeding 0.001, A-statistic = 0.68), we could not find a significant difference between platforms for effectiveness (p-value exceeding 0.05) although the effect size was found relevant (odds ratio = 0.41). Regarding satisfaction, most of our participants perceived the AWS as (i) having the best GUI to benefiting user interaction, (ii) the easiest platform to use, and (iii) the preferred cloud platform for creating VMs. Once confirmed by independent replications, our results suggest that AWS outperforms Azure regarding QiU. Therefore, cloud

Give Feedback

beginners might have a better cloud experience starting off their cloud projects by using AWS rather than Azure. In addition, our results may help to explain the AWS's cloud leadership.

Node.js Scalability Investigation in the Cloud

— *Jiapeng Zhu (University of New Brunswick) ; Panagiotis Patros (University of Waikato, University of New Brunswick) ; Kenneth Kent (University of New Brunswick) ; Michael Dawson (IBM Canada Ltd.) ;*

Short Abstract:

Node.js has gained popularity in cloud development due to its asynchronous, non-blocking and event-driven nature. However, scalability issues can limit the number of concurrent requests while achieving an acceptable level of performance. To the best of our knowledge, no cloud-based benchmarks or metrics focusing on Node.js scalability exist. This paper presents the design and implementation of Ibenchjs, a scalability-oriented benchmarking framework, and a set of sample test applications. We deploy Ibenchjs in a local and isolated cloud to collect and report scalability-related measurements and issues of Node.js as well as performance bottlenecks. Our findings include: 1) the scaling performance of the tested Node.js test applications was sub-linear; 2) no improvements were measured when more CPUs were added without modifying the number of Node.js instances; and 3) leveraging cloud scaling solutions significantly outperformed Node.js-module-based scaling.

Session Name: Machine Learning
Day: Tuesday, October 30th
Time: 8:30 - 10:00
Room: Primrose
Session Chair: Ayse Bener (Ryerson University)

A Context-Aware Machine Learning-based Approach

— *Nathalia Nascimento (Pontificia Universidade Católica do Rio de Janeiro (PUC-Rio)) ; Carlos Lucena (Pontificia Universidade Católica do Rio de Janeiro (PUC-Rio)) ; Paulo Alencar (University of Waterloo) ; Donald Cowan (University of Waterloo) ;*

Give Feedback

Short Abstract:

It is known that training a general and versatile Machine Learning (ML)-based model is more cost-effective than training several specialized ML-models for different operating contexts. However, as the volume of training information grows, the higher the probability of producing biased results. Learning bias is a critical problem for many applications, such as those related to healthcare scenarios, environmental monitoring and air traffic control. In this paper, we compare the use of a general model that was trained using all contexts against a system that is composed of a set of specialized models that was trained for each particular operating context. For this purpose, we propose a local learning approach based on context-awareness, which involves: (i) anticipating, analyzing and representing context changes; (ii) training and finding machine learning models to maximize a given scoring function for each operating context; (iii) storing trained ML-based models and associating them with corresponding operating contexts; and (iv) deploying a system that is able to select the best-fit ML-based model at runtime based on the context. To illustrate our proposed approach, we reproduce two experiments: one that uses a neural network regression-based model to perform predictions and another one that uses an evolutionary neural network-based approach to make decisions. For each application, we compare the results of the general model, which was trained based on all contexts, against the results of our proposed approach. We show that our context-aware approach can improve results by alleviating bias with different ML tasks.

Evaluating Music Mastering Quality Using Machine Learning

— *Mark Shtern (York University) ; Pedro Casas (York University) ; Vassilios Tzerpos (York University) ;*

Short Abstract:

Machine learning has been applied in a vast array of applications in the recent years, including several qualitative problems in the arts. However, in the world of music production, including mixing and mastering, most tasks are still performed by music professionals with decades of experience. Aspiring mastering engineers typically have to apprentice with professionals to learn their craft. Access to professionals is a scarce resource though, as they are typically very busy. In this paper, we present a method to evaluate the mastering quality of a piece of music automatically. We delegate the task of determining what we deem to be a subjectively well mastered song to professional mastering engineers. Using professionally mastered music, we derive datasets with varying degrees of deviation from the original music and train models to recognize the changes that have been

Give Feedback

made. This allows us to provide novice mastering engineers with an automatic rating of their work based on the magnitude of the deviation from the gold standard. We present experiments that demonstrate the accuracy of our approach, as well as a user study that shows how the results of our approach correlate to assessments made by human evaluators.

(Semi)Automatic Construction of Access-Controlled Web Data Services

— *Kalvin Eng (University of Alberta) ; Diego Serrano (University of Alberta) ; Eleni Stroulia (University of Alberta) ; Jacob Jaremko (University of Alberta) ;*

Short Abstract:

The widespread adoption of the Internet of Things (IoT) is producing an ever-increasing stream of data that can be mined by multiple stakeholders, in support of different objectives and tasks. In fact, we are witnessing the emergence of data marketplaces that aim to share this data and harness economic value out of these transactions. The advent of data-as-a-service (DaaS) represents a key integrator opportunity that allows for the management of data collections, while providing specific privacy policies to delegated agents. To support DaaS integrations, we develop a model-driven method for creating APIs to deliver DaaS. Our method supports data owners to: (1) automatically abstract the representation of relational database schemas into a visual model and map them to existing ontologies, (2) use the mappings in order to create different role-based access-control views of APIs, and (3) automatically generate API endpoints and their responses, based on these mappings. We develop a 'plug-and-play' prototype system for SQL databases to demonstrate this methodology and apply it to a use case of controlling data from a fitness monitoring application. Our aim is to enhance existing API creation methodologies that may be cumbersome by using semantics so that data can be easily shared and distributed.

Give Feedback

Papers

Think: Cognitive Computing, Big Data, Cloud, Security and Privacy

October 29 - 31

CASCON News

Download the CASCON 2018 Program Guide

See the CASCON 2018 Proceeding in ACM

Get certain Workshops materials

Call for Papers for CASCON 2019

Give Feedback

CASCON 2018 Proceeding can be found from **ACM**

Technical papers of CASCON 2018 should be original **research papers** and be at most 12 pages long in ACM format(style ACM_SigConf). CASCON 2018 **experience reports** share experience and lessons gained from the use of technology that can be useful to others and should be at most 12 pages in ACM format. CASCON 2018 **position papers** share vision, highlight technology gaps, or discuss emerging advances and should be at most 8 pages in ACM format. CASCON 2018 will not accept submissions that have been previously published, are in press, or have been submitted elsewhere. Accepted technical papers will be included in the conference proceedings published by CASCON and included in the ACM Digital Library.

Best Paper Awards

Two paper awards - **Best Paper and Best Student Paper** - will 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 student(s) must have primarily authored the paper, must have been a student(s) at the time the work was done, and must have carried out the work described; only the student author(s) receive this award.

Congratulations to the Winners!

The Best Paper Award goes to "Reducing Variability of Technically Related Software Systems in Large-Scale IT Landscapes".

→ Click to read more and share

The Best Student Paper Award is given to "Just-in-time Detection of Protection-Impacting Changes on WordPress and MediaWiki ".

→ Click to read more and share

Give Feedback

Most Influential Paper Award

Since 2010, CASCON has presented a **"Most Influential Paper" (MIP) Award** to a paper published a decade earlier at CASCON. The MIP Award recognizes lasting contributions and impact to theory and practice.

Congratulations to the Winners!

The Most Influential Paper from CASCON 2008 is awarded to "Is it a Bug or an Enhancement? A Text-based Approach to Classify Change Requests".

→ Click to read more and share

Give Feedback

Congratulations to the Authors!

Accepted Papers

This year we received 68 full papers and 23 position papers.
We accepted 23 full papers and 10 position papers.
Acceptance rate for full papers was 33.8% and for position papers was 43.5%.

Thank you for your submissions. See you all at CASCON 2018.

Full Papers | **Position Papers**

Session Name: Microservices
Day: Monday, October 29th
Time: 10:15 - 10:45
Room: Evergreen
Session Chair: Hausi Müller (University of Victoria)

Adaptation as a Service

— *Hamzeh Khazaei (University of Alberta) ; Alireza Ghanbari (Khatam-ol-Anbia (PBU) University) ; Marin Litoiu (York University) ;*

Short Abstract:

Current and emerging complex systems of many types including but not limited to big data systems, web-based systems, data centers and cloud infrastructure, social networks and the Internet of Things (IoT) have increasingly distributed and dynamic architecture that provide unprecedented flexibility in creating and supporting applications. However, such highly distributed architecture also increases the complexity of end-to-end management of such systems. Due to the sheer complexity, uncertainty and at the same time programmability of cloud environments, microservices and finally big data analytics, it is now required, and possible, to enable autonomic management in distributed systems in a dependable manner. In this paper, we argue that building autonomic management systems is a challenging task and requires its own set of expertise and knowledge. Therefore, in the light of current challenges, available enablers and recent successful stories, we propose the idea of moving from self-adaptation to ADaptation-as-a-Service (ADaaS).

Give Feedback

Uncertainty Quantification-as-a-Service

— *Malgorzata Zimon (IBM) ; Vadim Elisseev (IBM) ; Robert Sawko (IBM) ; Samuel Antão (IBM) ; Kirk Jordan (IBM) ;*

Short Abstract:

Uncertainty quantification (UQ), which enables non-destructive virtual testing, is the fast growing area of modern computational science. UQ methods are computationally intensive and require construction of complex work-flows, which rely on a number of different software components often coming from different projects. Therefore, there is a need for developing a portable and scalable UQ pipeline that will enable efficient stochastic modelling. Our paper introduces a strategy for UQ as a Service using high performance computing and hybrid cloud infrastructures and presents its application to a heat transfer study in nuclear reactors simulation and modelling of tsunami events.

Session Name: Adaptive Systems
Day: Monday, October 29th
Time: 10:15 - 10:45
Room: Jasmine
Session Chair: Jin Li (PointClickCare)

A DevOps Framework for Quality-Driven Self-Protection in Web Software Systems

— *Nasim Beigi Mohammadi (York University) ; Marin Litoiu (York University) ; Mahsa Emami-Taba (University of Waterloo) ; Ladan Tahvildari (University of Waterloo) ; Marios Fokaefs (Polytechnique Montreal) ; Ettore Merlo (Polytechnique Montreal) ; Vio Onut (IBM Canada Ltd.) ;*

Short Abstract:

Modern software is developed, deployed and operates continuously. At the same time, cyberattacks are on the rise. The continuity of development and operations and the constant threat of attacks requires novel approaches to identify, analyze and address potential security vulnerabilities. In this continuous and volatile execution environment, factors like security,

Give Feedback

performance, cost and functionality may not be able to be guaranteed in the same degree at the same time. In this work, we propose a DevOps framework for security adaptation that enables the development and operations teams to collaborate and address security vulnerabilities. The proposed framework spans across the different phases of software (development, operations, maintenance) and considers all other factors (performance, cost, functionality), when deciding for security adaptations. We demonstrate the approach on a prototype tool that shows how teams work together to tackle security concerns.

Ontology Driven Temporal Event Annotator mHealth Application Framework

— *Amente Bekele (Carleton University) ; Joe Samuel (Carleton University) ; Shermeen Nizami (Carleton University) ; Amna Basharat (National University of Computer & Emerging Sciences) ; Randy Giffen (IBM Canada Ltd.) ; James Green (Carleton University) ;*

Short Abstract:

We present an application (app) framework to facilitate the collection of gold standard temporal event annotations. These data will enable training and evaluation of machine learning algorithms for predicting events of clinical significance. Recording of such data using pen and paper can prove to be tedious and error-prone due to the variation in the types of events and the frequency of occurrence. To address this problem, we developed an mHealth application framework that presents an intuitive and configurable user interface for annotating a timeline with events. The presented Temporal Event Annotator (TEA) app framework supports dynamically building a customized application inclusive of events, event categories, and study attributes based on the design input of a specific study. This is accomplished by presenting a terminology schema for the hierarchical definition of event types and an additional user interface (UI) schema to support UI-specific attributes. We describe the framework architecture independent of specific technology implementations. We also describe specific instantiations of the framework that we used to develop and evaluate apps for three different use cases: 1) patient monitoring in the Neonatal Intensive Care Unit (NICU), 2) estimating patient stress levels during immersive rehabilitation therapy, and 3) quantifying the patient experience during emergency neonatal transport. The TEA framework provides a reliable and intuitive solution for temporal event annotation that accounts for the unique experimental requirements of each study.

Give Feedback

Session Name: Compiler Development
Day: Monday, October 29th
Time: 10:15 - 10:45
Room: Primrose
Session Chair: Andrew Craik (IBM)

A Survey of Ahead-of-Time Technologies in Dynamic Language Environments

— *Mark Thom (University of New Brunswick) ; Gerhard Dueck (University of New Brunswick) ; Kenneth Kent (University of New Brunswick) ; Daryl Maier (IBM Canada Ltd.) ;*

Short Abstract:

Eclipse OMR is an open source collection of robust, reuseable components for the construction of production-ready compilers. Great progress has been made on JITBuilder, OMR's simplified interface to the compiler technology for building JIT compilers, but work on the planned interface for AOT compilation has only just begun. In this survey, we identify desirable characteristics for the design of OMR's AOT by examining how several prominent open source compilers implement AOT. We conclude by discussing the advantages and disadvantages of the implementations seen, and how they might inform the final design of the OMR AOT component.

All Timescale Window Co-occurrence

— *Yumeng Liu (University of Rochester) ; Daniel Busaba (University of Rochester) ; Chen Ding (University of Rochester) ; Daniel Gildea (University of Rochester) ;*

Short Abstract:

Trace analysis is a common problem in system optimization and data analytics. This paper presents new efficient algorithms for window co-occurrence analysis, which is to find how likely two events will occur together in time windows of different lengths. The new solution requires a linear time preprocessing step, after which, it only takes logarithmic space and constant time to compute co-occurrence of a data pair in windows of any given length. One potential use of the new analysis is to reduce the asymptotic cost in affinity-based memory layout.

Give Feedback

Session Name: Blockchain
Day: Tuesday, October 30th
Time: 10:15 - 10:45
Room: Jasmine
Session Chair: Biruk Habtemariam (IBM)

Challenges and solutions on architecting Blockchain Systems

— *Gregory Fournier (Polytechnique Montreal) ; Fabio Petrillo (University of Quebec at Chicoutimi) ;*

Short Abstract:

Despite the fact that companies are gravitating more and more towards the use of blockchains in their systems, it is clear that the blockchains is no silver bullet. Many challenges such as scalability issues and frustrating trade-offs most notable in public decentralized blockchain systems are currently holding back blockchain's huge potential. In this paper we conduct a Systematic Literature Review in order to explore the current challenges of blockchain while presenting possible solutions to each of these challenges. We come to the conclusion that current challenges can be summarized in three categories: Scalability issues, security issues and a choice of consensus protocol. We also briefly discuss the use of blockchain in current systems, concluding that while blockchains current immaturity makes it hard to recommend for most projects, blockchains in their current state could be used in the Internet of Things.

Powering Software Sustainability with Blockchain

— *Omar Badreddin (University of Texas) ;*

Short Abstract:

Software sustainability is a systematic challenge that impacts broad segments of software systems. Software codebases must evolve overtime to address changing contexts and adapt to the flux in middlewares and platforms. In the process, it accumulates arbitrary complexities and its maintenance becomes progressively difficult. Current sustainability approaches

Give Feedback

focus on the symptoms and tend to be reactive in nature, and ignore the fundamental incentive structures that drive decision-making processes. Moreover, contemporary approaches are insensitive to the uniqueness of each software project context and operate on the assumption that sustainability measurements are universally applicable to the majority of software systems.This paper introduces a fundamentally novel peer-driven approach to managing software sustainability. The methodology ensures that software teams can define their own sustainability measures that adequately address the unique context of their project and its priorities. These measures are dynamically defined by the project peers to ensure their applicability as the project context evolves. Finally, the paper introduces Susereum, a blockchain platform that materializes the methodology and establishes novel incentive structures to systematically promote software sustainability throughout the project lifecycle.

Session Name: Healthcare
Day: Tuesday, October 30th
Time: 10:15 - 10:45
Room: Primrose
Session Chair: Fred Popowich (Simon Fraser University)

Ischemic Brain Stroke Detection using EEG signals

— *Arooj Ahmed Qureshi (McMaster University) ; Canxiu Zhang (McMaster University) ; Rong Zheng (McMaster University) ; Ahmed Elmeligi (HiNT) ;*

Short Abstract:

Stroke is the second leading cause of death in the United States of America. 87% of all strokes are ischemic stroke, which is mainly caused by the blockage of small blood vessels around the brain. Magnetic resonance imaging (MRI) provides the gold standard for accurate diagnosis of ischemic strokes, but it is both time-consuming and unsuitable for 24/7 monitoring. In this paper, we propose an ischemic stroke detection method through the multi-domain analysis of EEG brain signal from wearable EEG devices and machine learning. Using 40 healthy and 40 patients' data, we find that Multi- Layered Perceptron (MLP) and Bootstrap models (Extra-Tree and Decision-Tree) can achieve test accuracy of 95% with an area under the ROC curve 0.85.

Give Feedback

Predictive Analytics in Healthcare: Epileptic Seizure Recognition

— *Ashok Bhowmick (Ryerson University) ; Tamer Abdou (Arish University) ; Ayse Bener (Ryerson University) ;*

Short Abstract:

Introduction Clinical applications of electroencephalography (EEG) span a very broad range of diagnostic conditions. Epileptic seizure is the fourth most common neurological disorder in that. Related Work There has been considerable progress in clinical understanding of epilepsy, however many aspects of seizure prevention are still a mystery. Predictive modeling of EEG can provide significant value addition to substantiate the diagnosis of epilepsy. Methodology Machine learning algorithms are applied to predict the probability of epileptic seizure using an open source multi-class dataset. Results and Discussion Comparing the F-score from different classifiers, it is found that XGBoost gives the best performance in binary classification and Random Forest provides the best performance in multinomial classification. Conclusion Our results show that it is possible to predict epileptic seizure with significant accuracy from non-epileptic parameters using a suitable machine learning algorithm. We also observe that binary classification methods have higher prediction accuracy.

Give Feedback

Workshops

Think: Cognitive Computing, Big Data, Cloud, Security and Privacy

October 29 - 31

CASCON News

Download the CASCON 2018 Program Guide

See the CASCON 2018 Proceeding in ACM

Get certain Workshops materials

Call for Papers for CASCON 2019

Give Feedback

Materials for certain CASCON 2018 workshops are now available from **HERE**

Workshops at CASCON 2018 will provide a forum to present, discuss, and debate issues, problems, ideas, technology gaps, work-in-progress, and/or directions. The format of a workshop may include position papers, expert panels, hands-on exercises, and discussions. All submitted workshop proposals require long abstracts of 1500 words maximum, typically including the abstract,

rationale, technical/research scope, organizers, workshop format, and expected outcomes. Long abstracts of accepted workshops will be included in the conference proceedings published by CASCON and included in the ACM Digital Library.

Congratulations to the Workshop Chairs!

Accepted Workshops

Thank you for your submissions. See you all at CASCON 2018.

Monday PM workshops run from 03:15 to 05:15 PM

— Building AI models using IBM Watson Studio *(HandsOn, in Conf Center 2)*

Chair(s):	Sarah Packowski , Wendy Switzer
Theme:	Cognitive Computing
Room:	Conf Center 2
Format:	HandsOn
Level:	Beginner
Prereq:	1. You MUST bring your own laptops for the hands-on exercise 2. Set up Watson Studio and Watson Knowledge Catalog on IBM Cloud: https://dataplatform.cloud.ibm.com/registration/stepone
Description:	In this workshop you will train AI models that process different types of sample input data, including: structured data (tabular) and unstructured data (images and sound.) Then you will see how to integrate those models into fun, sample apps that solve everyday challenges. With our help, you will create models in different ways using a variety of tools in IBM Watson Studio. About Watson Studio: IBM Watson Studio provides a range of tools to help you train AI models - from graphical tools that guide you, step-by-step, in choosing machine learning algorithms to notebooks where you can construct complex

Give Feedback

neural networks by hand. Watson Studio simplifies AI development, whether you are looking for tools to make learning AI easier, tools to rapidly prototype your AI inventions, or a cost-effective, powerful platform for AI research or enterprise AI solutions. See more: https://medium.com/ibm-watson/introducing-ibm-watson-studio-e93638f0bb47

View Workshop Detail

— Build a cognitive serverless Slack app with IBM Cloud Functions & IBM Watson API *(HandsOn, in Holly-Butternut)*

Chair(s):	Serjik Dikaleh , Eric Charpentier , John Liu , Neil DeLima , Vince Yuen
Theme:	Cloud Computing
Room:	Holly-Butternut
Format:	HandsOn
Level:	Intermediate
Prereq:	- Some JavaScript and Node.js knowledge - Sign up for an IBM Cloud and a Slack account
Description:	Slack is an easy use collaboration tool to serve as a digital communication hub in many companies and teams. Based on Apache OpenWhisk, IBM Cloud Functions is a functions-as-a-service (FaaS) programming platform for developing lightweight code that scalably executes on demand. This workshop will teach the audience to build a Slack app by implementing several serverless IBM cloud functions and integrating them into slack channels through Slack Events API. The application will also leverage IBM Watson APIs to have a Slack chatbot to chat with users to provide cognitive services within a demonstrated knowledge domain.

View Workshop Detail

— Build better APIs with the next generation of API testing and monitoring *(HandsOn, in Orchid)*

Chair(s):	ivy ho , JJ Tang , JISOO LEE , Amirali Jafarian , Peter El-koreh
Theme:	Data and Analytics
Room:	Orchid
Format:	HandsOn
Level:	Beginner

Give Feedback

Prereq: None. Nice to have some knowledge of what APIs are.

Description: May need own laptop with https://www.ibm.com/cloud/api-connect/api-test signed up before joining the workshop. The proliferation of APIs across all facets of life continued to explode and grow. The quality of the APIs and the data retrieval become a critical factor. In this workshop, we will walk you through a new no-code way of validating the API quality. How you can validate the API payload accuracy. How you can get new insights into API data from real business use-cases in different sectors. You will witness the innovation of the API test and monitor approach in this workshop .

View Workshop Detail

— **Come learn how to deploy Open Liberty applications using Docker, Kubernetes, Helm and MicroProfile!** *(HandsOn, in Violet)*

Chair(s): Arthur De Magalhaes , Leo Christy Jesuraj
Theme: Cloud Computing
Room: Violet
Format: HandsOn
Level: Beginner
Prereq:
Description: Application modernization is in everyone's mind - but what environment do you migrate your legacy application into? Can that environment also host your new cloud-native applications?

In this hands-on lab you'll learn about how to leverage Open Liberty's Docker container from Open Source to package your applications (legacy or cloud-native) into a container and store them in a private, on-premises Docker registry.

You will then use IBM's Open Source Helm charts to deploy these applications (plus a database!) into Kubernetes using production-grade settings such as auto-scaling and health monitoring.

Lastly, you will see how MicroProfile OpenAPI can help your containerized microservices expose their REST APIs and enable an internal API economy between legacy and cloud-native applications.

The environment for this lab will be IBM Cloud Private - a production-ready kubernetes platform.

View Workshop Detail

Give Feedback

— Best Practices and Lessons Learned in Microservices *(Panel, in Evergreen)*

Chair(s): Julia Rubin , Yingying Wang , Harshavardhan Kadiyala , John Steinbacher , Tony Erwin
Theme: Cloud Computing
Room: Evergreen
Format: Panel
Level: Beginner
Prereq:
Description: Microservice-based architecture is an approach to developing a single application as a suite of independent services. The services run in separate processes and communicate with each other via lightweight language-agnostic protocols, such as HTTP REST. The services are split following business capabilities; each service has a fully automated pipeline and is independently deployable.

Microservices aim at shortening the development lifecycle while improving the quality, availability, and scalability of applications at runtime. From the development perspective, cutting one big application into small independent pieces reinforces the component abstraction, and makes it easier for the system to maintain clear boundaries between components. At runtime, microservices can be individually scaled by adding more instances of those microservices that experience increasing traffic.

Due to these advantages, microservice-based architectures are now becoming increasingly popular in industry. Examples of companies that have been using microservices include Amazon, Netflix, IBM, Uber, LinkedIn, Groupon, and eBay.

Yet, adopting microservice-based architectures and implementing it “right” is no a trivial endeavour. Just “jumping on the microservices trend” and expecting that the transition itself, together with the adoption of advanced technology, such as Docker and Kubernetes, will allow companies to achieve significant improvements is a false belief.

In this workshop, we intend to explore best practices, lessons learned, and technical challenges practitioners face when adopting and implementing microservices. These include considerations for identifying the right service granularity and topology, issues related to synchronization and constancy, security of microservices, performance debugging, efficient monitoring and troubleshooting, and more.

Our goal is to gather researchers and practitioners interested in exchanging ideas on the topic. For practitioners, the workshop will provide venue to learn from each other, borrow successful ideas, and avoid common mistakes. For researchers, a description of current practices and challenges practitioners face can inspire novel software

Give Feedback

engineering methods and techniques.

The workshop will be structured as a series panel discussions and invited talks by participants from industry and academia.

View Workshop Detail

— Data-driven medicine: promise and challenges *(Panel, in Primrose)*

Chair(s): TOMAS TOKAR , Igor Jurisica
Theme: Data and Analytics
Room: Primrose
Format: Panel
Level: Beginner
Prereq:
Description: In the last few years, use of machine learning (ML) and artificial intelligence (AI) in the sector of healthcare started to gain broader acceptance. These technologies may revolutionize medicine by improving diagnostic accuracy and increasing therapeutic efficiency. The successful application of these technologies requires a constant circulation of a large amounts of data from patients, through healthcare professionals, to scientists and software developers. The data are not only required for the initial training and testing of the algorithms, they are also essential for monitoring algorithms' performance once deployed in a clinical practice. Nowadays medical data may include range of modalities. These include genetic profiles, medical imaging records, data from wearable devices, clinical findings and various socioeconomic characteristics of patients. This poses several technical and ethical challenges, addressing which requires introduction of a novel technologies and development of a new healthcare policies. This is only possible through broad communication between the medical professionals, computer science experts and information privacy specialists.

View Workshop Detail

— 13th Workshop on Challenges For Parallel Computing *(Speakers, in Conf Center 1)*

Chair(s): Jeeva Paudel , wael yehia , Jeremy Bradbury
Theme: Systems

Give Feedback

Room: Conf Center 1
Format: Speakers
Level: Beginner
Prereq:
Description: Parallel computing has expanded significantly over the past decade and now includes the development of applications for multi-core systems, distributed systems and heterogeneous systems. The goals of this workshop are to bring together different groups from the parallel community (application developers, language developers, compiler and tools developers, system architects and academic researchers) to explore the current challenges that parallel computing faces and present ideas on how to deal with these challenges.

View Workshop Detail

— The 3rd International Workshop on Dew Computing *(Speakers, in Elm2)*

Chair(s): Yingwei Wang , Karolj Skala
Theme: Cloud Computing
Room: Elm2
Format: Speakers
Level: Beginner
Prereq: None
Description: DEWCOM is an annual international workshop on dew computing. The first one, DEWCOM 2016, was held in Charlottetown, Canada. The second one, DEWCOM 2017, was held in Opatija, Croatia. The third one, DEWCOM 2018, is generously sponsored by IBM Centre for Advanced Studies and CASCON 2018 and will be held together with CASCON 2018. The details of DEWCOM 2018 can be found in http://www.dewcomputing.org/index.php/dewcom-2018/.

Dew computing is a new post-cloud computing model appeared in 2015. While cloud computing uses centralized servers to provide various services, dew computing uses on-premises computers to provide decentralized, cloud-friendly, and collaborative micro services to end-users.

Dew computing is an on-premises computer software-hardware organization paradigm in the cloud computing environment, which does not contradict with cloud computing, does not replace cloud computing, but it is complementary to cloud computing. The key features of dew computing are that on-premises computers provide functionality independent of cloud services and they also collaborate with cloud services. Briefly speaking, dew computing is a better way of using local computers in the age of cloud computing.

Give Feedback

In this workshop, research progress in dew computing will be presented. Ideas and future directions will be discussed.

This workshop will have 5 sessions. The schedule of this workshop is: Session 1: Oct. 29, 3:15 – 5:15pm. Session 2: Oct. 30, 8:30 – 10:00am. Session 3: Oct. 30, 10:15am. – 12:00pm. Session 4: Oct.30, 1:00 – 3:00pm. Session 5: 3:15 – 5:15pm. Here we introduce the content of the first session.

This session will include “Dew Computing Tutorial” and a presentation: “Dewblock: A Blockchain System Based on Dew Computing.”
In the tutorial, we will focus on the following aspects: What is dew computing? What are the features of dew computing? Where can dew computing be applied to?

The presentation is about an application of dew computing to the blockchain technology. Blockchain is great and has huge potentials, but the size of a blockchain is always increasing. This will eventually cause problems for the use of blockchains. In this presentation, a new kind of blockchain system based on dew computing, Dewblock, will be introduced. The major feature of this new blockchain is that the data size of a client is very small and the features of a full node are still kept. This approach opens the door for the blockchain technology to be widely used in personal computers and mobile devices.

View Workshop Detail

— IBM Academic Skills Academy - Syllabi and Other Things *(Speakers, in Jasmine)*

Chair(s):	Stephen Perelgut , Dennis Buttera , Colette Lacroix , Lila Adamec
Theme:	Other
Room:	Jasmine
Format:	Speakers
Level:	Intermediate
Prereq:	This workshop is intended primarily for academic faculty and administrators
Description:	IBM has initiated a program to "Teach the Teacher" how to use the latest technologies. This session details 4 of the most requested topics, giving the complete syllabus for the 40hr course as well as highlights from the available badges for: Explorer, Mastery and Instructor.

Attendees will leave with a clear idea of what they can expect from the Skills Academy program and how they can

Give Feedback

learn materials to pass on to others.

Topics will include:
- Business Intelligence Analytics
- Mobile Application Development and IoT
- Blockchain and Design Thinking for Blockchain
- Quantum Computing

View Workshop Detail

Give Feedback

Workshops

Think: Cognitive Computing, Big Data, Cloud, Security and Privacy

October 29 - 31

CASCON News

Download the CASCON 2018 Program Guide

See the CASCON 2018 Proceeding in ACM

Get certain Workshops materials

Call for Papers for CASCON 2019

Give Feedback

Materials for certain CASCON 2018 workshops are now available from **HERE**

Workshops at CASCON 2018 will provide a forum to present, discuss, and debate issues, problems, ideas, technology gaps, work-in-progress, and/or directions. The format of a workshop may include position papers, expert panels, hands-on exercises, and discussions. All submitted workshop proposals require long abstracts of 1500 words maximum, typically including the abstract, rationale, technical/research scope, organizers, workshop format, and expected outcomes. Long abstracts of accepted workshops will be included in the conference proceedings published by CASCON and included in the ACM Digital Library.

Congratulations to the Workshop Chairs!

Accepted Workshops

Thank you for your submissions. See you all at CASCON 2018.

Tuesday PM workshops run from 03:15 to 05:15 PM

— Blockchain Fundamentals and Development Platforms *(HandsOn, in Conf Center 2)*

Chair(s):	Omar Badreddin
Theme:	Security
Room:	Conf Center 2
Format:	HandsOn
Level:	Beginner
Prereq:	No required prerequisites. All required software is available in the cloud. Any web browser will be sufficient to perform the development tasks. Users MUST bring their own laptops for the hands-on exercise.
Description:	Blockchain is an emerging computing and development platform. It is a new paradigm that aims at empowering peers and eliminate the need for central authenticating authority. In essence, blockchain has introduced a novel level of distributed sovereignty. Crypto currency is one prominent outcome of this new computing paradigm that has gained

Give Feedback

broad attention. However, Blockchain paradigm has demonstrated broader potential impacts in many disciplines including secure software engineering, supply chain, banking, and peer to peer commerce.

This half day workshop will give participants a brief background on the technology along with hands on practice on development using some prominent open source blockchain platforms. The hands-on exercises will be led by experienced Blockchain developer.

The first talk in the workshop will introduce the fundamental Blockchain concepts. The next talk will provide an overview of current and emerging impacts of Blockchain technologies covering many disciplines and industries, including financial industry, supply chains, authentication and security, as well as the recent emergence of blockchain based distributed social networks.

After the first two talks, participants will engage with guided hands-on exercises to develop a basic blockchain application. Participants will be given the required software and will also be made available online.

By the end of the workshop, participants are expected to have gained in-depth understanding of the emerging Blockchain technology and its applications. Participants will also gain knowledge and skills on existing blockchain development platforms. Therefore, this workshop is ideal for researchers and practitioners alike who are considering using Blockchain in their research or work, as well as middle technical managers who want to gain understanding on how Blockchains and the emergence of distributed sovereignty may impact their current line of businesses. The workshop is useful for educators who wish to introduce Blockchain in their undergraduate and graduate courses.

View Workshop Detail

— Modernize digital applications with Microservices management using the Istio service mesh *(HandsOn, in Holly-Butternut)*

Chair(s):	Ozair Sheikh , Serjik Dikaleh , Dharmesh Mistry , Darren Pape , Chris Felix
Theme:	Cloud Computing
Room:	Holly-Butternut
Format:	HandsOn
Level:	Intermediate
Prereq:	- Intermediate Kubernetes knowledge - A free trial IBM Cloud (Bluemix) account - Minimal familiarity with Linux command line
Description:	Digital solutions are being built on modernized enterprise platforms deployed on cloud infrastructure and managed

Give Feedback

using container platforms. Foundational infrastructure capabilities such as load balancing and routing, previously available as software are now being provided as part of the underlying cloud platform. When designing your next generation architecture, its integral to understand the capabilities available from the cloud platform versus acquiring / developing it with software. For example, load balancing and automatic scaling are features that are built-into container orchestration platforms such as Kubernetes; therefore, you should not expect your applications to develop these capabilities, rather write your applications in a manner that allows you to embrace the container platform.

These key application design principals are based on API / Microservices architecture, where business functions are packaged and deployed within containers and communicate with each other using API interfaces. As the number of microservices grow, the need to manage the interactions and provide key runtime capabilities becomes a critical requirement for success ... lets explore why the service mesh is the right architecture for microservices-based applications.

A service mesh is an infrastructure layer for controlling container (ie microservice-to-microservice) traffic in microservices-based applications. Each container (ie microservice) is deployed together with a separate “sidecar” proxy, which interacts with a “control plane”, enforcing access between microservices. The service mesh provides a clear boundary between runtime operations and microservices functionality. It standardize the runtime operations using a declarative approach, so you can write policies to enforce runtime behaviour without developing any code. For example, lets explore the circuit breaker pattern. This pattern helps prevent failure for your entire application when a single service / component is unresponsive. Netflix Hystrix is a popular library used within Java applications to provide circuit breaker functionality. The challenge with using a shared library is that it gets embedded in code and becomes difficult to manage when code changes need to be made; instead, using an out-of-process proxy (ie sidecar) allows your microservice to add circuit breaker capability without modifying your application.

Istio is an open community project that implements the service mesh architecture. It is built on top of Kubernetes and provides an additional runtime layer that adds scalability, routing, A/B testing and more. It allows you to inject a “side proxy” into an existing kubernetes pod without changing your application, reducing the friction for adoption. You automatically gain the benefits of telemetry, security and circuit breaking without writing code or modifying configuration.

In this workshop, you will learn how to deploy an single-page application (SPA) built with API/Microservices design principals into the Istio service mesh. You will get hands-on experience in configuring Istio-based policies to manage microservices interaction (ie service routing) and improve infrastructure resiliency (ie circuit breaker) without writing application code. The overall application resiliency is enhanced when you deploy your application within the Istio service mesh.

View Workshop Detail

Give Feedback

— Hands-On: Easy Microservices Application Development with Microclimate *(HandsOn, in Orchid)*

Chair(s):	Elson Yuen , Eric Peters , Rajiv Senthilnathan , Maysun Jamil Faisal , Steven Hung
Theme:	Cloud Computing
Room:	Orchid
Format:	HandsOn
Level:	Beginner
Prereq:	Basic knowledge on Java and JavaScript is recommended.
Description:	Microclimate is a brand-new, cloud native development environment that offers a complete, end-to-end development experience for Microservices. Since Microclimate has been designed with a focus on containerization, it can run anywhere from your local laptop, to an IBM Cloud private cluster.

With Microclimate, you can create or import Java, Node.js, or Swift applications into the development environment, and using any editor of your choosing, you can quickly start development on your application in a containerized environment. Through a process called Rapid Iteration, Microclimate will quickly detect any changes that occur in your project and determine the minimal and best course of action to update your application. From there, using our integrated DevOps pipeline, you can deploy your application with Jenkins to a live ICP cluster. With these features, Microclimate offers a fully featured development experience that many other environments don't offer today.

During the hands-on workshop, we will give you an introduction to Microclimate, starting from product installation to write Microservices applications to run on Microclimate in a Docker environment. You will get hands on experiences to create new applications and import existing applications into Microclimate. For developers, a crucial part of the development cycle is the ability to quickly develop and test applications changes on a running application. The develop-deploy-test-repeat cycle must be as short as possible in order to prevent lost developer productivity due to deployment downtime. You will be given the opportunity to experience this rapid iterative development support by developing Java and JavaScript applications in this workshop.

Finally, during the workshop we will introduce the integrated DevOps pipeline functions provided that allows you get into production fast with a preconfigured DevOps pipeline and deploy application to IBM Cloud Private (ICP). We will also show you the diagnostic services that helps you to do problem determination in production.

View Workshop Detail

Give Feedback

— IBM Voice Agent with Watson *(HandsOn, in Violet)*

Chair(s): Alice Yeung , Rick Chen , Philip Kurowski , Trevor Crawford , Meswan Bhaugeerutty
Theme: Cognitive Computing
Room: Violet
Format: HandsOn
Level: Beginner
Prereq: Attendees need to create trial Twilio and IBM Cloud accounts during the workshop. The accounts are free and no credit card information is needed.
Description: Cognitive chat bots are changing the way businesses are interacting with their customers. Whether embedded in a web page, or talking to users via a mobile application, Watson powered cognitive bots can resolve queries quickly and efficiently.

To better leverage this powerful technology, businesses can use Voice Agent with Watson on IBM Cloud to quickly build Watson powered chat bots (voice agents) and connect them to the telephone network.

Through connecting Speech to Text, Text to Speech, and Watson Assistant, voice agents can identify what a caller is saying and to respond back in real time. As Watson Assistant can process natural language, it is able to converse with callers using complete sentences, helping to improve the experience for callers.

Voice agents can act as a self-service solution to solve the most common calls entering call centers. By answering the main volume of calls with voice agents, businesses can resolve calls more effectively. This results in a reduction to the volume of calls sent to human agents.

View Workshop Detail

— Third Annual Workshop on Data-Driven Knowledge Mobilization *(Panel, in Jasmine)*

Chair(s): Kelly Lyons , Eleni Stroulia , Marcellus Mindel
Theme: Data and Analytics
Room: Jasmine
Format: Panel
Level: Beginner
Prereq: None

Give Feedback

Description: Knowledge mobilization and translation describes the process of moving knowledge from research and development (R&D) labs into environments where it can be put to use. There is increasing interest in understanding mechanisms for knowledge mobilization, specifically with respect to academia and industry collaborations. At the same time, the number of available datasets and accessible analysis tools is growing.

Building on the discussions and results presented at previous workshops, the third annual workshop on data-driven knowledge mobilization will bring together researchers, students, and industry partners to present results and discuss challenges associated the analysis of datasets associated with knowledge mobilization. In order to understand the processes of knowledge mobilization, we need access to certain datasets and specific analyses techniques. We will present details of curated datasets and analysis techniques that support analysis of individuals and resources, expertise and work activities, and work outputs and innovations.

The goals of the workshop are to bring participants together to share results and outcomes and to discuss challenges and future directions. In this workshop, we will report on research underway as part of a Strategic Partnership Project on Data-Driven Knowledge Mobilization, Translation, and Innovation. The Strategic Partnership Project is in its final year. Since the start of the project, several people have participated in the research including six investigators, two postdoctoral fellows, eight Ph.D. students, five masters students, and fourteen industrial and partner scientists. One of the goals of the project is to produce a repository of datasets and analysis tools. The theme of this 3rd annual workshop focuses on tools to enable the analysis of datasets that will help in understanding the processes of knowledge mobilization.

There will be five presentations by researchers and students involved in the project followed by a panel discussion.

View Workshop Detail

— The 10th CASCON Workshop on Cloud Computing *(Speakers, in Conf Center 1)*

Chair(s): Marin Litoiu , Joe Wigglesworth
Theme: Cloud Computing
Room: Conf Center 1
Format: Speakers
Level: Intermediate
Prereq:
Description: The goal of the workshop was to bring together researchers and practitioners from government, industry and academia to present and share the best practices and research agendas at the intersection of Cloud Computing and

Give Feedback

Internet of Things: development, deployment, runtime management, quality of services and runtime models. We particularly focused on several main topics: cloud requirements for Internet of Things, deployment and adaptive runtime management, cognitive capabilities, security and privacy. DevOps plays an important role in the IoT and cloud ecosystems, providing the mechanisms that enable agile development and operations and it was a topic of the workshop. Application domains such as smart buildings and smart cities were illustrated.

This half-day workshop consisted in presentations and a panel. The presentations were structured along the main themes of the workshop. To encourage discussion and provide a more open discussion and perspective, we included a panel where industry and academic experts presented their visions and answered questions from the audience.

View Workshop Detail

— CASCON Workshop on Developing Big Data Applications and Services - BDAS 2018 *(Speakers, in Elm2)*

Chair(s):	Darlan Arruda , Nazim H. Madhavji , Colin Taylor
Theme:	Other
Room:	Elm2
Format:	Speakers
Level:	Intermediate
Prereq:	To make the most of this workshop, participants must have a working (general) knowledge of software engineering, software development, and Big Data.
Description:	Nice to have: Industry experience in developing Big Data Applications and Services. 1. Background Research from Gartner (2015) indicates that, in 2017, 60% of Big Data projects failed or did not provide the expected benefits. However, in November 2017, Nick Heudecker, a Gartner analyst, posted in his twitter account that they were too conservative. The Big Data project failure rate is now close to 85%. The reasons are not only related to technology itself. It is a mix of environmental, technological and managerial problems. Some of the reasons for Big Data projects failure are: At the project level: missing link to business objectives, lacking big data skills, relying too much on the data, failing to convince executives, and poor planning; At the technical level: Rapid technology changes, difficulty in selecting Big Data technologies to address the systems and project requirements, complex integration between new and old systems, computation of intensive analytics, and the necessity of high scalability, availability and reliability, to name a few. Further, our previous study has shown that there is approximately a 80:20 split in the industry focus in favour of “algorithms for analytics” and “infrastructure”, thereby shortchanging the aspects of creating and evolving “applications” and “services” concerned with Big Data.

Give Feedback

2. Importance
The emerging data on project challenges or failure should be of immense concern to the Big Data software community. It calls for meeting of the minds to deliberate about, and share experiences concerning the development of Big Data applications and services. Both industry and academia representation are needed to cut through the barriers facing the community today. Working in isolation may prolong the pain and agony of challenges faced in the Big Data software community. In turn, the society at large is deprived of the potential benefits of Big Data applications and services. The proposed workshop comes at a critical juncture in the fast-emerging field of Big Data applications and services development. The workshop aims to be a catalyst in the movement on Big Data applications building and services creation. It will form a platform for participants, from both practice and research, to deliberate on, and achieve a deeper understanding of, the different activities, methods, tools, processes, system artefacts, constraints, conditions, etc., involved in Big Data projects.

3. Purpose and Interest
Given the described importance of the proposed workshop, the purpose of the workshop includes: (1) sharing Big Data project experiences among the participants, and identifying challenges in the design, implementation, deployment, and evolution of Big Data applications and services; (2) fostering a Big Data community of researchers and practitioners focused on applications and services, and (3) compiling an agenda for future research. All application domains are of interest in this workshop. In the workshop, we shall identify, debate about, and discuss solutions to, the barriers challenging the development, deployment, evolution, and success of Big Data Applications and Services.

View Workshop Detail

— The Best of IBM Innovation: Advancements through Overcoming Technological Uncertainties *(Speakers, in Evergreen)*

Chair(s):	Jerrold Landau , Perry Fuller
Theme:	Data and Analytics
Room:	Evergreen
Format:	Speakers
Level:	Beginner
Prereq:	None necessary. Beginners and experienced people welcome.
Description:	There are many factors that spur innovation in the field of technology. One such factor is governmental support. The Canadian government provides such support through the SR&ED (Scientific Research and Experimental Development) tax credit program. Not surprisingly, IBM has availed itself of this tax credit for many years. The fundamentals of

Give Feedback

SR&ED are based on three pillars: a) identifying a technological uncertainty, b) overcoming the uncertainty through a scientific experimental process, c) leading to an advancement in technology. The advancement is often defined as the acquisition of new knowledge in the domain. It should be noted that the advancement need not be incorporated into a product, and indeed need not necessarily be successful in the classical business definition of the term. One advances technology just as well by proving that something is not feasible than by proving something is feasible. This definition of advancements in technology, while nuanced and subject to governmental SR&ED program regulations, can be applied in a most general sense to many fields of scientific innovation. It has been noted that both the IBM CAS organization and IBM's participation in the SR&ED program promote innovation through an exploration of the proverbial 'bleeding edge' of technology. In this workshop, we will provide an overview of the technological criteria for participation in the SR&ED program, and present SR&ED as an indicator of technical vitality. This will be followed by presentations from several IBM Lab teams highlighting their SR&ED claims over the past several years. It is expected that the audience will come away with a new perspective on scientific innovation as applied to the field of software development.

View Workshop Detail

— Large-Scale Multilevel Streaming Data Analytics *(Speakers, in Primrose)*

Chair(s): Farhana Zulkernine , Haruna Isah
Theme: Data and Analytics
Room: Primrose
Format: Speakers
Level: Beginner
Prereq:
Description: Motivation and Justification:

There is a monumental shift happening in how data powers organizational and business operations. This shift is about moving away from traditional batch data analytics to real-time and hybrid data analytics involving both static and continuous data to avoid delay in generating insights and storing massive amount of streaming data. A good number of analytics systems currently utilize stream processing without storing the data to quickly ingest, analyze and to correlate information as it arrives from thousands of real-time sources (devices, sensors, and applications). Such systems often provide real time dashboards and critical alerts, and therefore, are required to be fast, efficient, effective, scalable, and reliable.

In most cases stream processing is followed by batch processing for deeper analytical processing. Modern streaming

Give Feedback

analytic systems, therefore, try to unify batch and streaming analytics into a seamless data processing pipeline. A general architecture of a large-scale multilevel analytics system consists of (i) an ingestion mechanism at the front-end, (ii) streaming and batch data processing engines for data transformation, scoring, modelling of historical data, and real-time prediction, (iii) data storage units for persisting, indexing, searching, and knowledge management, (iv) resource management unit for the coordination of distributed compute and storage resources, and (v) visualization units to present results and knowledge for decision support.

Some of the deeper analytics of streaming data requires longer execution time and can choke the data processing pipeline. The stream plus batch analytics solves that problem. However, in our progression towards the Internet of Things (IoT), we will face serious computational and storage challenges in such an approach. Innovative solutions are needed to selectively store streaming data, enable near real time micro batch processing, and perform multi-level in-memory analytics.
Large-scale multilevel analytics on a unified platform is increasingly gaining attention in the industry as it can potentially enhance business and operational decision making. However, it faces the following challenges, a) implementing an efficient front-end for ingestion and integration of massive data streams across the globe, b) combining streaming and in-memory data analytics, c) developing a knowledge management strategy to store, manage and link big data and distributed knowledge, and d) other challenges including cluster management, knowledge representation, and visualization. The above challenges make the development of methods, algorithms, and infrastructures for multilevel streaming analytics a challenging but interesting research problem.

Goals and Outcomes:

This workshop aims to provide a forum for researchers and industry practitioners to discuss new ideas and share their experiences in the areas of streaming data analytics. Participants will present their work on topics including methods, models, algorithms, infrastructures, quality issues, applications, and open problems for large-scale streaming data analytics. The workshop can serve as a guide for organizations and individuals planning to implement a real-time data stream processing and multilevel data analytics framework.

Workshop Structure:

The half-day workshop will feature invited talks by experts, practitioners, researchers, and industry partners working on massive streaming analytics research. There will be a time for discussion after each presentation to instigate the audiences to share their comments and views and ask questions to the speaker.

View Workshop Detail

Give Feedback

Workshops

Think: Cognitive Computing, Big Data, Cloud, Security and Privacy

October 29 - 31

CASCON News

Download the CASCON 2018 Program Guide

See the CASCON 2018 Proceeding in ACM

Get certain Workshops materials

Call for Papers for CASCON 2019

Give Feedback

Materials for certain CASCON 2018 workshops are now available from **HERE**

Workshops at CASCON 2018 will provide a forum to present, discuss, and debate issues, problems, ideas, technology gaps, work-in-progress, and/or directions. The format of a workshop may include position papers, expert panels, hands-on exercises, and discussions. All submitted workshop proposals require long abstracts of 1500 words maximum, typically including the abstract,

rationale, technical/research scope, organizers, workshop format, and expected outcomes. Long abstracts of accepted workshops will be included in the conference proceedings published by CASCON and included in the ACM Digital Library.

Congratulations to the Workshop Chairs!

Accepted Workshops

Thank you for your submissions. See you all at CASCON 2018.

Monday PM	Tuesday PM	**Wednesday AM**	Wednesday PM	Wednesday Full Day

Wednesday AM workshops run from 08:30 to 10:45 AM

— **Building Microservices in a Cloud-Native World using Eclipse MicroProfile and Open Liberty** *(HandsOn, in Orchid)*

Chair(s): Eveline Cai , Gilbert Kwan , YK Chang , Panagiotis Roubatsis
Theme: Cloud Computing
Room: Orchid
Format: HandsOn
Level: Beginner
Prereq:
Description: Eclipse MicroProfile is a set of open technologies to optimize enterprise Java for a microservices architecture. Open Liberty is the open source foundation of WebSphere Liberty, IBM's strategic Java application server for a cloud-native world. Come and learn how you can easily build microservices with Eclipse MicroProfile and Open Liberty. Get your hands dirty with various aspects of building cloud-native applications, from foundation for RESTful services to what you need to manage many microservices and what you need for reliable operations.

View Workshop Detail

Give Feedback

— Refine, restructure and make sense of data visually, using IBM Watson Studio *(HandsOn, in Violet)*

Chair(s): Serjik Dikaleh , Darren Pape , Dharmesh Mistry , Chris Felix , Ozair Sheikh
Theme: Cognitive Computing
Room: Violet
Format: HandsOn
Level: Beginner
Prereq: Must sign up for IBM Cloud account prior to the workshop (trial account/free tier is fine) - https://www.ibm.com/cloud/
Optional basic coding knowledge
Description: More than ever before, larger and more comprehensive data sets are being made publicly available on the internet. You can find data on all sorts of topic such as housing prices, sports data, wine reviews, weather, movies, TV shows, gun violence and anything you can think of. How does one make sense of all this data, and can you combine different data sets to get new insights for your needs?

In this workshop, we will begin by talking about where you can find open data sets and show some examples of how they have been used to gain insights. We will then take some sample data sets and explore it through IBM Watson Studio. We will further the workshop by creating visualizations of the data using both open source programming concepts and using tools available in IBM Watson Studio.

View Workshop Detail

— 2nd Workshop on DevOps and Software Analytics for Continuous Engineering and Improvement *(Speakers, in Jasmine)*

Chair(s): KONSTANTINOS KONTOGIANNIS , Chris Brealey , Alberto Giammaria , Brian Countryman , Marios-Stavros Grigoriou
Theme: Systems
Room: Jasmine
Format: Speakers
Level: Intermediate
Prereq: Knowledge on software development process and software engineering principles
Knowledge on DevOps tools and frameworks
Experience in software development and its life-cycle
Description: A key issue that emerges in the software engineering community is how to provide efficient DevOps tools and processes that facilitate continuous delivery and improvement, particularly in and for Cloud based environments

Give Feedback

where continuous delivery at speed with high quality can be crucial to business success.
This workshop aims to bring experts from industry and academia to discuss and debate the latest trends related to the design of frameworks that support DevOps practices of complex systems which are developed and evolved within a “Measure-Analyze-Assess-Act” loop. Such frameworks utilize software repositories, software analytics, process analytics, the quantification of technical debt as a failure risk predictor, and the system’s run-time behavior to dynamically assess deploy/no-deploy choices and achieve continuous deployment.

View Workshop Detail

Give Feedback

Workshops

Think: Cognitive Computing, Big Data, Cloud, Security and Privacy

October 29 - 31

CASCON News

Download the CASCON 2018 Program Guide

See the CASCON 2018 Proceeding in ACM

Get certain Workshops materials

Call for Papers for CASCON 2019

Give Feedback

Materials for certain CASCON 2018 workshops are now available from **HERE**

Workshops at CASCON 2018 will provide a forum to present, discuss, and debate issues, problems, ideas, technology gaps, work-in-progress, and/or directions. The format of a workshop may include position papers, expert panels, hands-on exercises, and discussions. All submitted workshop proposals require long abstracts of 1500 words maximum, typically including the abstract, rationale, technical/research scope, organizers, workshop format, and expected outcomes. Long abstracts of accepted workshops will be included in the conference proceedings published by CASCON and included in the ACM Digital Library.

Congratulations to the Workshop Chairs!

Accepted Workshops

Thank you for your submissions. See you all at CASCON 2018.

Wednesday PM workshops run from 02:15 to 05:30 PM

— IBM Security Guardium Analyzer Bootcamp *(HandsOn, in Orchid)*

Chair(s):	Devan Shah , Larry Lindsay , Josue Diaz , Sagi Shechter , Andy Becher
Theme:	Security
Room:	Orchid
Format:	HandsOn
Level:	Beginner
Prereq:	IBM ID account already created
Description:	General Data Protection Regulation (GDPR) requires organizations to implement adequate controls to protect personal and sensitive personal. A critical step in that journey is impact assessment; understanding where data (subject to GDPR) is located and how vulnerable it is. In this session, we will introduce IBM Security Guardium Analyzer, a new SaaS offering, which intends to help organizations easily complete the impact assessment for their databases. You will receive hands-on experience on quickly utilizing the Guardium Analyzer solution to locate GDPR

Give Feedback

data and determining risk on existing databases.

View Workshop Detail

— Deriving Client Insights in the Financial Sector *(HandsOn, in Violet)*

Chair(s):	Diane Reynolds , DAVID DCOSTA , David Xie , Seacy Zhen
Theme:	Data and Analytics
Room:	Violet
Format:	HandsOn
Level:	Intermediate
Prereq:	- solid knowledge of coding in Python - exposure to wealth management / financial sector - willingness to participate in team/group activities
Description:	Join us in this hands-on workshop to create your own reports and dashboards to support a financial advisor in completing key elements of his/her day-to-day activities. We'll look at data requirements, experiment hands-on with the data, clean it, load it to a data-science-friendly environment, run some standard models and then extend those models in different ways. Finally, we'll bring together the results in a user-friendly way. If you've wondered about how to operationalize your machine learning algorithms, want to get deeper into data science as a financial-sector participant or are interested in IBM's ecosystem for machine learning innovation, this is the right workshop for you!

View Workshop Detail

— iCity - Big Data and Visualization Urban Transportation Strategies *(Speakers, in Jasmine)*

Chair(s):	Sara Diamond
Theme:	Data and Analytics
Room:	Jasmine
Format:	Speakers
Level:	Beginner
Prereq:	There are no prerequisites for the workshop.
Description:	Providing efficient, cost-effective, sustainable transportation networks and services is a major challenge for cities

Give Feedback

around the world – not only for individual cities, but for connectivity between cities. High quality transportation services, notably well-designed transit hubs within comprehensive networks are fundamental prerequisites for effective cities and spur economic, social and cultural inclusion, development and growth. Transportation strategies must be at the heart of smart city strategies. The melding of machine learning, simulations, predictive analytics and design create capacity and connectivity that will help policy and makers gain insight into complex decision-making processes and support evidence-based decision making. Solving transportation and transit challenges requires integrating transdisciplinary knowledge, including computer science, engineering into city planning.

View Workshop Detail

Give Feedback

Workshops

Think: Cognitive Computing, Big Data, Cloud, Security and Privacy

October 29 - 31

CASCON News

Download the CASCON 2018 Program Guide

See the CASCON 2018 Proceeding in ACM

Get certain Workshops materials

Call for Papers for CASCON 2019

Give Feedback

Materials for certain CASCON 2018 workshops are now available from **HERE**

Workshops at CASCON 2018 will provide a forum to present, discuss, and debate issues, problems, ideas, technology gaps, work-in-progress, and/or directions. The format of a workshop may include position papers, expert panels, hands-on exercises, and discussions. All submitted workshop proposals require long abstracts of 1500 words maximum, typically including the abstract, rationale, technical/research scope, organizers, workshop format, and expected outcomes. Long abstracts of accepted workshops will be included in the conference proceedings published by CASCON and included in the ACM Digital Library.

Congratulations to the Workshop Chairs!

Accepted Workshops

Thank you for your submissions. See you all at CASCON 2018.

Monday PM	Tuesday PM	Wednesday AM	Wednesday PM	**Wednesday Full Day**

Wednesday Full-Day workshops run from 08:30 to 10:45 AM, and then continue from 02:15 to 05:30 PM.

— Practical Machine Learning with Python on DSX *(HandsOn, in Conf Center 2)*

Chair(s):	Shaikh Quader , Mark Ryan , Eric Dong
Theme:	Data and Analytics
Room:	Conf Center 2
Format:	HandsOn
Level:	Beginner
Prereq:	* Coding experience in any programming language * Users MUST bring their own laptops for the hands-on exercise * IBM Cloud id: https://console.bluemix.net/registration/
Description:	In this FULL-DAY workshop, after a brief introduction to Machine Learning, we'll take the students through the hands-on exercise of building a Machine Learning model from the scratch. They will learn and code in different phases of Machine Learning pipeline, including acquiring, cleaning, and exploring data; building and evaluating ML model.

Give Feedback

Finally, we'll tell them how to build a discipline for continuous learning of ML and applying this to solving real problems.

View Workshop Detail

— Introduction to the IBM Q experience and Quantum Computing *(HandsOn, in Holly-Butternut)*

Chair(s):	Mehdi Bozzo-Rey , Robert Loredo
Theme:	Other
Room:	Holly-Butternut
Format:	HandsOn
Level:	Beginner
Prereq:	It is assumed that participants will: - bring their own laptops - have the anaconda python distribution installed (www.anaconda.com) - have an IBM Q experience account (quantumexperience.ng.bluemix.net)
Description:	IBM's work in quantum computing started in the 1970s with the birth of quantum information theory and the first conference on Physics of Computation was co-hosted by IBM and MIT in 1981. The quantum foundations time, where quantum computing was the exclusive domain of scientists and theoreticians is past history. We are now in a quantum readiness phase where education, algorithm development and use case identification that may lead to demonstrations of quantum advantage are key. In 2016 IBM made quantum computing capable devices available for the public at no cost in the cloud and then launched the IBM Q network in 2017. Programming a quantum computer is quite different from “classical programming” and is done with a circuit made of quantum gates that execute the quantum algorithm. For devices with a small number of qbits, a GUI driven interface can be used to build the circuit, making quantum programming available and viable for even high schoolers. For devices with more qubits, or more complex circuits a programming framework is needed. In this hands-on workshop, we will review the basics of quantum computing, go through account creation on the IBM Q experience and basic use of the Composer (GUI driven interface), installation of the Open Source Qiskit and Qiskit AQUA frameworks, and how to execute simple quantum circuits on either a local quantum simulator or on a real quantum device that operates at a temperature colder than outer space.

View Workshop Detail

Give Feedback

— Compiler-Driven Performance Workshop *(Speakers, in Conf Center 1)*

Chair(s): Gennady Pekhimenko , Ettore Tiotto
Theme: Systems
Room: Conf Center 1
Format: Speakers
Level: Intermediate
Prereq:
Description: The workshop has a particular focus on (but not limited to):
Innovative compiler analysis, transformation, and optimization techniques
Languages, compilers, and optimization techniques for multicore processors and other parallel architectures
Compiling for streaming or heterogeneous hardware
Dynamic compilation for high-performance and real-time environments
Compilation, optimization, and analysis for dynamic languages
Compilation techniques for reducing power
Program safety
Whole system optimization and analysis
Tools and infrastructure for compiler research

View Workshop Detail

— Distributed Ledgers and Blockchain: Concepts and Applications *(Speakers, in Evergreen)*

Chair(s): Asic Chen , Arno Jacobsen
Theme: Other
Room: Evergreen
Format: Speakers
Level: Beginner
Prereq: General computer science and computer engineering (information systems background as taught in Bachelor's curricula)
Description: Blockchain has been, without doubt, one of the hottest topics in technology in recent years. As is in the case of most "buzzwords", most people, even software professionals, have no more than a surface comprehension of the technology. We hope that through this workshop, we can demystify blockchain and distributed ledgers, giving the

Give Feedback

attendees a working understanding while expanding on use cases far beyond the most popular: cryptocurrency. During this workshop, we will provide a tutorial-style introduction to various distributed ledgers and blockchain technologies. We will be focusing on first principles and algorithms while identifying emerging blockchain use cases.

View Workshop Detail

— 2nd Workshop on Advances in Open Runtime Technology for Cloud Environments *(Speakers, in Primrose)*

Chair(s): Daryl Maier , Kenneth Kent
Theme: Cloud Computing
Room: Primrose
Format: Speakers
Level: Intermediate
Prereq:
Description: Modern language runtimes are complex, dynamic environments that involve a myriad of components that must work cooperatively to achieve the functional and performance requirements of a given language. Typical core runtime technologies include dynamic just-in-time compilers for performance, garbage collection for heap management, platform abstraction for ease of portability to different hardware and operating system environments, developer tooling for diagnosis and tuning of the various components, and interoperability between different language environments.

Cloud services such as IBM Cloud, Microsoft Azure, or Amazon Web Services (AWS) are increasingly becoming the environments where applications are developed and deployed, data is stored, and businesses are run. Many of the features that define a cloud (e.g., resiliency, elasticity, consistency, security) are realized through runtime technologies. Clouds are polyglot environments, and therefore advances in cloud development are directly driven by innovation in runtime technologies. However, cloud environments pose unique, often conflicting demands on runtime systems that are often less of a concern in isolated systems. Throughput performance (how fast is my application?), density (how many instances of my application can I run simultaneously in my provisioned environment?), startup performance (how quickly can I launch a new instance of my application?), and language interoperability (how can my JavaScript code efficiently call a function in a Python module?) are all important considerations that require innovation to solve effectively.

The goal of this workshop is to bring together research, industry, and development communities to share and discuss innovations, challenges, and research across a broad set of open-source runtime technologies (such as Eclipse OMR, LLVM, Eclipse OpenJ9, Node.js) for cloud environments. The focus on open technology solutions rather than

Give Feedback

proprietary is key as it allows for greater collaboration amongst individuals, communities, researchers, and companies through shared learning on common technology.

View Workshop Detail

Give Feedback

Expo

Think: Cognitive Computing, Big Data, Cloud, Security and Privacy

October 29 - 31

CASCON News

Download the CASCON 2018 Program Guide

See the CASCON 2018 Proceeding in ACM

Get certain Workshops materials

Call for Papers for CASCON 2019

Give Feedback

The Technology Expo at CASCON 2018 will provide an excellent opportunity to highlight your research results and generate interest in upcoming product and product areas. Exhibits on emerging technologies are encouraged. The format of an exhibit may be a demonstration, poster presentation, or a short talk. A booth is provided at no cost to the exhibitors. Each demonstration booth will be

set up with power and wireless internet access. Exhibitors are responsible for bringing the necessary equipment such as machines, laptops, monitors, power cables, and security locks; and any printed materials including posters.

Accepted exhibits

Congratulations to the exhibit chairs and thank you for your submissions. See you all at CASCON 2018.

Monday | Tuesday | Wednesday

Exhibit Sessions - 11:00 to 13:00 and 17:30 to 19:00

— *B1:* Document Signing on Blockchain

Booth Number:	B1
Title:	Document Signing on Blockchain
Chair(s):	Asic Chen , Arno Jacobsen
Format:	Demo
Description:	Our project provides a fast, secure, and decentralized solution for document signing based on blockchain technology. Document endorsements are immutably recorded and reliably time-stamped on the blockchain, easily verifiable by third parties that neither need to authenticate the signatories, nor need to access the content of the document. User identities are verified without the need for a centralized certificate authority. Our prototype is built on Hyperledger Fabric with an Angular front end.
Industry Company:	IBM
Web Site:	http://www.msrg.org/
Research Team:	**University of Toronto:** Asic Chen Arno Jacobsen Yuxi (Eric) Zheng Chen (Jackie) Gao Yuxuan (Eric) Zheng Bogdan Scaunasu

Give Feedback

— *B2:* PowerAI - Democratizing AI for All Research Communities

Booth Number:	B2
Title:	PowerAI - Democratizing AI for All Research Communities
Chair(s):	Marija Mijalkovic
Format:	Demo
Description:	IBM's PowerAI platform provides a unique solution to support, enhance and progress the ground breaking research pursuits of diverse academic institutions and faculties. The solution is underpinned by the most powerful and innovative infrastructure for the development of artificial intelligence and layers on data science tools and utilities to enable both skilled computer science and mathematics experts as well as those with 'domain expertise' who are knowledgeable in their respective fields.
Business unit	IBM Cognitive Systems
Product Name:	AC922, PowerAI Enterprise
Research Team:	**IBM:** Marija Mijalkovic

— *B3:* Data Science 1 - Competency Model

Booth Number:	B3
Title:	Data Science 1 - Competency Model
Chair(s):	Marcellus Mindel , Dennis Buttera , Jennifer Collins , Ana Echeverri
Format:	Poster
Description:	In the first of three exhibits on Data Science Education, we will discuss applications of a Competency Model for Data Science originally developed by Ana Echeverri, IBM Data Science Growth Strategies Lead, to support the nontraditional education paths of new collar jobs. The model may also prove useful for postsecondary institutions that use competency matrices to manage the coverage of data and analytics skills in academic programs.
Industry Company:	IBM
Research Team:	**IBM Advanced Studies:** Marcellus Mindel Dennis Buttera Jennifer Collins **IBM Data Science Growth Strategies:** Ana Echeverri

Give Feedback

— *B4:* Performance and Security Enhanced Containerized Analytics for IoT Systems

Booth Number:	B4
Title:	Performance and Security Enhanced Containerized Analytics for IoT Systems
Chair(s):	Simon Nadeau , Nathaniel Holeksa , Marios Fokaefs
Format:	Poster
Description:	Our research delves into the pseudo dichotomy that is to be secure or to be performant, in a cloud based IoT architecture. We explore the two in the context of a sensor network in smart buildings, a novel technology that hopes to improve our environmental footprint. In our work, we deploy Spark streaming jobs in Docker containers and explore how to effectively, and automatically, scale them using Docker Compose pursuant to the CPU and memory usage of the standard and the security analytics.
Industry Company:	IBM
Research Team:	**Polytechnique Montreal:** Simon Nadeau Nathaniel Holeksa Marios Fokaefs **York University:** Brian Ramprasad

— *B5:* Randomized Offloading Algorithms for Mobile Cloud Computing

Booth Number:	B5
Title:	Randomized Offloading Algorithms for Mobile Cloud Computing
Chair(s):	Haleh Shahzad , Majid Risman Kar Zadeh
Format:	Poster
Description:	Computational offloading can improve the energy efficiency of mobile devices, by executing some tasks of a mobile application in the cloud. In my research, a set of new algorithms are presented to find the best decision for each task. These algorithms iteratively improve an offloading decision vector, trying to find the optimal one that can minimize the total energy consumption of the mobile device when the application is running while satisfying the delay requirement of the application.
Industry Company:	IBM
Research Team:	**McMaster University:** Haleh Shahzad

Give Feedback

Rogers Communications: Majid Risman Kar Zadeh

— *B6:* Exploring Challenges In Adoption of Continuous Delivery

Booth Number:	B6
Title:	Exploring Challenges In Adoption of Continuous Delivery
Chair(s):	Trevor Rae
Format:	Poster
Description:	We look at the technical and cultural problems that exist in regards to adopting continuous delivery/integration into a company's infrastructure.
Industry Company:	IBM
Product Name:	IBM Cloud Continuous Delivery Service
Research Team:	**University Of Victoria:** Trevor Rae Dr.Damian

— *B7:* Container Technologies! Which One Best Suits Eclipse IoT Deployments?

Booth Number:	B7
Title:	Container Technologies! Which One Best Suits Eclipse IoT Deployments?
Chair(s):	Mohab Aly , Foutse Khomh
Format:	Poster
Description:	The increasing apex in the IoT paradigm requires specific platforms to allow different components to communicate. Our research dives into assessing the performance implications of the Open-Sourced Eclipse IoT, including Hono, to provide guidelines while deploying them in production. In our work, we setup loads in container technologies, such as Kubernetes and OpenShift, and explore how to effectively scale them out using EnMasse while analyzing the CPU usage, memory consumption and network I/O.
Industry Company:	IBM
Research Team:	**Polytechnique Montréal:** Mohab Aly Foutse Khomh

— *B8:* Fall as an irregular pattern in IoT streaming data

Give Feedback

Booth Number: B8
Title: Fall as an irregular pattern in IoT streaming data
Chair(s): Sazia Mahfuz
Format: Poster
Description: Detecting patterns in real time streaming data has been an interesting and challenging data analytics problem. In this work, we address the problem of learning an irregular human activity pattern, fall, from streaming IoT data using wearable sensors. We present a deep neural network model for detecting fall based on accelerometer data using an online dataset "MobiAct" by Vavoulas et al. The initial model was developed using IBM Watson studio and then later transferred and deployed on IBM Cloud.
Industry Company: IBM
Research Team: **Queen's University:** Sazia Mahfuz Dr. Farhana Zulkernine
IBM Canada: Peter Nicholls

— *B9:* DevOps Round-trip Engineering: Traceability from Dev to Ops and Back

Booth Number: B9
Title: DevOps Round-trip Engineering: Traceability from Dev to Ops and Back
Chair(s): Miguel Jimenez
Format: Demo
Description: DevOps engineers follow an iterative and incremental process to develop Deployment and Configuration specifications. Such process is likely to involve manual modifications to the running environment. Failing to appropriately update the specifications leads to technical debt. In this demo, we present a two-way Continuous Integration framework that automatically updates Terraform templates when the target OpenStack deployment changes, enabling bi-directional traceability of the modifications.
Industry Company: IBM
Web Site: https://github.com/RigiResearch/jachinte-DevOps2018-evaluation
Research Team: **University of Victoria:** Miguel Jimenez Hausi Müller
Universidad Icesi: Norha M. Villegas Gabriel Tamura
IBM Canada: Ian Watts Joe Wigglesworth

Give Feedback

— *B10:* Installing features to the Open Liberty kernel with Maven and Gradle

Booth Number:	B10
Title:	Installing features to the Open Liberty kernel with Maven and Gradle
Chair(s):	Eric Lau
Format:	Poster
Description:	Open Liberty is a lightweight and modular server runtime for Java developers. It has features on Maven Central that can be installed to its kernel. We showcase the use of these features, the steps to add them into Maven and Gradle builds, and the benefits of doing so for developers. We explain how they provide relevant dependencies for compiling applications which simplifies the development process. We also provide high level details of the components used to make this solution possible.
Business unit	IBM Hybrid Cloud
Product Name:	Open Liberty
Web Site:	https://openliberty.io/
Research Team:	**IBM:** Eric Lau Alvin So Patrick Tiu

— *B11:* A Hierarchical Rule-based Security Management System for Data-Intensive Applications

Booth Number:	B11
Title:	A Hierarchical Rule-based Security Management System for Data-Intensive Applications
Chair(s):	Marin Litoiu , Yar Rouf , Vio Onut
Format:	Demo
Description:	We propose a platform to support the SecDevOps framework, a hierarchical distributed architecture for security control that uses a Business Rules Engine. The Business Rules Engine simplifies security rules by allowing the teams to write them at an operational level rather than at the network level, which would require specialized knowledge.
Industry Company:	IBM
Product Name:	QRadar, Cloud Automation Manager

Give Feedback

Research Team: **York University:** Yar Rouf Marin Litoiu Joydeep Mukherjee Vio Onut
IBM: Vio Onut

— *B12:* Continuous Delivery with Open Toolchains

Booth Number: B12
Title: Continuous Delivery with Open Toolchains
Chair(s): Chris Brealey , Curtis dEntremont , Christophe ELEK , Joel Cayne
Format: Demo
Description: Continuous Delivery (n) - 1. The subtle science and exact art of releasing high quality, bite-sized changes to software at high velocity without sacrificing control or the ability to pivot with the needs of our clients. - 2. a feature of the IBM Cloud, built on the Open Toolchain platform, that delivers industry leading development tools and tool integrations as a service. Drop by to see toolchains in action, and to get the creative juices flowing on ways to evolve them!
Business unit IBM Watson & Cloud Platform
Product Name: IBM Cloud Continuous Delivery
Web Site: https://bluemix.net/devops
Research Team: **IBM DevOps Services:** Colleen Anderson Chris Brealey Joel Cayne Curtis D'Entremont Christophe Elek Mike Melick Philippe Mulet Simon Kotwicz Herniczek Tom Schmidt Herman Singh and many more on Dave Thomson's tribe

— *B13:* Ahead-Of-Time Compilation of WebAssembly Using Eclipse OMR

Booth Number: B13
Title: Ahead-Of-Time Compilation of WebAssembly Using Eclipse OMR
Chair(s): Mark Thom , Gerhard Dueck , Daryl Maier , Kenneth Kent
Format: Poster
Description: WebAssembly is an emerging assembly language that runs on the web. It has nearly native performance, and is already the compilation target of several programming languages. There is an ongoing project to

Give Feedback

create a JIT compiler for WebAssembly using Eclipse OMR's JITBuilder. The poster will describe our attempted contribution of an AOT compiler implemented with OMR. We hope it will expose any disparity between OMR's current support for AOT and what its completed AOT module should be capable of.

Industry Company: IBM
Project Number 1036
Product Name: Eclipse OMR
Web Site: http://www.unb.ca/research/casatlantic/index.html
Research Team: **University of New Brunswick:** Mark Thom Gerhard Dueck Kenneth Kent
IBM: Daryl Maier

— *B14:* Requirements-based Framework for Evaluating Trust in Service-oriented Systems

Booth Number: B14
Title: Requirements-based Framework for Evaluating Trust in Service-oriented Systems
Chair(s): KONSTANTINOS KONTOGIANNIS , KOSTAS TSIOUNIS
Format: Poster
Description: - We define trust as the degree of expectation that a service, resource, or system, will satisfy its predefined functional and non-functional requirements. - Users define their requirement expectations sought by a service or a resource (i.e. their goals) as expressions. - Fuzzy goal model reasoners, the TidalTrust algorithm, and information obtain from the system are used to compute the final trust a user associates that the service will be achieving the goals set in a given context.
Industry Company: IBM
Research Team: **Western University:** KOSTAS KONTOGIANNIS KOSTAS TSIOUNIS HAO JIANG
IBM Canada: JOE WIGGLESWORTH

— *B15:* Tracking Unhandled Rejections in Node.js

Booth Number: B15
Title: Tracking Unhandled Rejections in Node.js

Give Feedback

Chair(s): Maxim Uzun , Kenneth B. Kent , Michael Dawson
Format: Poster
Description: Node.js is an open source cross-platform runtime environment for executing server-side JavaScript. The framework is based on Google Chrome's V8 and is limited in regard to post-mortem diagnostics and debugging for Promises; a design pattern for managing asynchronous code. We propose and experimentally evaluate an implementation of Node.js that aims to improve troubleshooting for unhandled Promise rejections. The goal is to produce accurate core dumps and track unhandled rejections.
Industry Company: IBM
Project Number 1002
Product Name: IBM SDK for Node.js
Web Site: http://www.unb.ca/research/casatlantic/index.html

— *B16:* A semi-automated Framework for Migrating Web applications from SQL to NoSQL database

Booth Number: B16
Title: A semi-automated Framework for Migrating Web applications from SQL to NoSQL database
Chair(s): Rahma Al Mahruqi
Format: Poster
Description: We propose a semi-automated approach to migrate highly dynamic SQL-based web applications to ones that uses NoSQL database. There are two parts to this approach: the migration of schema and data, and the migration of the actual application code with embedded queries. Our approach provides contributions to migrating and optimizing the embedded SQL queries to interact with the new database system and changing the application code to use the translated queries.
Industry Company: IBM
Research Team: **Queen's University:** Rahma Al Mahruqi Manar Alalfi Thomas Dean

— *B17:* Feature engineering in Big Data for the detection of information systems misuse

Booth Number: B17

Give Feedback

Title: Feature engineering in Big Data for the detection of information systems misuse
Chair(s): Eduardo Lopez , Kamran Sartipi
Format: Poster
Description: This cybersecurity research project explores the Apache Spark-enabled feature engineering activities — conceptualization and generation— in machine learning detection of an insider's threat. Our experiment uses a very large, 1.6 billion records dataset of information systems logs, collected for more than 12,000 users over a period of 58 days. We test the feature set using classification models, validating feature relevance and significance using Receiver Operating Characteristics curves.
Industry Company: IBM
Research Team: **McMaster University,East Carolina University:** Eduardo Lopez (IS at McMaster) Kamran Sartipi (CS at ECU and IS at McMaster

— *B18:* Composing Web Services Using a Multi-Agent Framework

Booth Number: B18
Title: Composing Web Services Using a Multi-Agent Framework
Chair(s): Yu Zhao
Format: Poster
Description: Software agents are autonomous entities that can proactively activate web services, and autonomously react upon the results to determine the next steps. However, programming agents are complex due to the the proactive, autonomous and reactive characteristics of agents. To reduce the workload of developers when programming agents, we propose an easy-to-understand semi-natural language syntax to program agents for service composition.
Industry Company: IBM
Project Number 920
Product Name: Building Intelligent Personal Agents
Research Team: **Queen's University:** Yu Zhao Daniel Alencar da Costa Ying Zou
IBM: Joanna Ng

Give Feedback

— *B19:* Detection of Feature Interaction in Dynamic Scripting Languages

Booth Number:	B19
Title:	Detection of Feature Interaction in Dynamic Scripting Languages
Chair(s):	Omar Al Harthi
Format:	Poster
Description:	Our research indicates that duplicate names, from published plugins, can be a potential cause of conflict due to overlapping. Such feature interaction may create a challenge for smooth compatibility of JavaScript plugin. To detect duplicate names we need first to understand how the objects and properties are created, used and changed in JavaScript.This work aims to detect such conflict using static analysis approach.
Industry Company:	IBM
Research Team:	**Queens University:** Omar Al Harthi Manar Alalfi Thomas Dean

— *B20:* Enabling a highly-scalable, hybrid cloud-based microservices architecture

Booth Number:	B20
Title:	Enabling a highly-scalable, hybrid cloud-based microservices architecture
Chair(s):	Anthony Kwan , Jonathon Wong
Format:	Demo
Description:	Traditional microservice scaling methods perform either horizontal or vertical scaling exclusively. When used in combination, however, these methods offer complementary benefits and compensate for each other's deficiencies. To leverage the high availability of horizontal scaling and the fine-grained resource control of vertical scaling, we developed two novel hybrid autoscaling algorithms and benchmarked them against Google's popular Kubernetes horizontal autoscaling algorithm.
Industry Company:	IBM
Product Name:	IBM Blueworks
Research Team:	**University of Toronto:** Anthony Kwan, Jonathon Wong, Hans-Arno Jacob **IBM:** Allen Chan, Suzette Samoojh, Vinod Muthuswam

Give Feedback

— *B21:* The Power of Accelerators for Cognitive Workloads

Booth Number:	B21
Title:	The Power of Accelerators for Cognitive Workloads
Chair(s):	Mel Bakhshi , Theresa Xu , Neil Graham
Format:	Poster
Description:	Moore's Law is no longer operative. Yet, cognitive workloads demand ever-more computational resources to process ever-increasing volumes of data. What options does he industry have to move forward? In this poster session, we'll show how accelerators such as GPUs, FPGAs and high-bandwidth storage and network can help to address these challenges and show how POWER systems are uniquely designed to take advantage of these emerging technologies.
Product Name:	Power 9 Systems
Research Team:	**IBM Power Systems Performance:** Neil Graham Charlie Wang Bill Phu Douglas Gibbs Theresa Xu Mel Bakhshi

— *B22:* Predicting Cross-Sell Opportunities in Financial Services Using Deep Learning

Booth Number:	B22
Title:	Predicting Cross-Sell Opportunities in Financial Services Using Deep Learning
Chair(s):	Chux Ejieh , Diane Reynolds
Format:	Poster
Description:	In this exhibition, we will explore the design and implementation of the Product Propensity model (part of the IBM Client Insights for Wealth Management SaaS) and share how it impacts our clients. Cross-sell refers to the practice of selling an additional product to an existing client. Ideally, marketing campaigns in this domain are more targeted as they leverage information the business has about its customers.
Industry Company:	IBM
Product Name:	Client Insights
Web Site:	https://www.ibm.com/us-en/marketplace/client-insight-for-wealth-management
Research Team:	**IBM Canada:** Chux Ejieh - Data Scientist Diane Reynolds - Chief Data Scientist

Give Feedback

— *B23:* Attention models for gene promoter sequence characterization

Booth Number:	B23
Title:	Attention models for gene promoter sequence characterization
Chair(s):	Raul Ivan Perez Martell , Alison Ziesel , Alex Thomo , Ulrike Stege
Format:	Poster
Description:	Augmented recurrent neural networks are well-suited to the analysis of eukaryotic promoter sequences. DNA sequences, like other biological sequences, are related to other sequential problems such as natural language processing (NLP). In this work we will investigate the efficacy of attention-augmented RNN networks on molecular biology data including DNA sequence, ChIP, and protein binding assays. We also investigate different biological data types and their relevance for a well-trained network.
Industry Company:	IBM
Research Team:	**University of Victoria:** R. I. Perez Martell A. Ziesel A. Thomo U. Stege

— *B24:* Using AP-TED to Detect Phishing Attack Variations

Booth Number:	B24
Title:	Using AP-TED to Detect Phishing Attack Variations
Chair(s):	Sophie Le Page
Format:	Poster
Description:	Detection of phishing attack variations is accordingly currently used as a prevention tool against phishing. We evaluate here the feasibility of applying Pawlik and Augsten's recent implementation of Tree Edit Distance (AP-TED) calculations as a way to compare DOMs and identify similar phishing attack instances. We also compare this tree method with an existing method that uses the distance between tag vectors to quantify similarity between phishing sites.
Industry Company:	IBM
Research Team:	**University of Ottawa:** Sophie Le Page

Give Feedback

Expo Reception / Short Talks / Expo Presentations

Short Talk Sessions starts from 17:30 (5 mins each). Presentations order as follows:

— *Monday B1 (17:30):* **Document Signing on Blockchain**

Exhibit Session:	Monday B1
Title:	Document Signing on Blockchain
Chair(s):	Asic Chen , Arno Jacobsen
Format:	Demo
Description:	Our project provides a fast, secure, and decentralized solution for document signing based on blockchain technology. Document endorsements are immutably recorded and reliably time-stamped on the blockchain, easily verifiable by third parties that neither need to authenticate the signatories, nor need to access the content of the document. User identities are verified without the need for a centralized certificate authority. Our prototype is built on Hyperledger Fabric with an Angular front end.
Industry Company:	IBM
Web Site:	http://www.msrg.org/
Research Team:	**University of Toronto:** Asic Chen Arno Jacobsen Yuxi (Eric) Zheng Chen (Jackie) Gao Yuxuan (Eric) Zheng Bogdan Scaunasu

— *Monday B15 (17:35):* **Tracking Unhandled Rejections in Node.js**

Exhibit Session:	Monday B15
Title:	Tracking Unhandled Rejections in Node.js
Chair(s):	Maxim Uzun , Kenneth B. Kent , Michael Dawson
Format:	Poster
Description:	Node.js is an open source cross-platform runtime environment for executing server-side JavaScript. The framework is based on Google Chrome’s V8 and is limited in regard to post-mortem diagnostics and debugging for Promises; a design pattern for managing asynchronous code. We propose and

Give Feedback

	experimentally evaluate an implementation of Node.js that aims to improve troubleshooting for unhandled Promise rejections. The goal is to produce accurate core dumps and track unhandled rejections.
Industry Company:	IBM
Project Number	1002
Product Name:	IBM SDK for Node.js
Web Site:	http://www.unb.ca/research/casatlantic/index.html

— *Monday B16 (17:40):* **A semi-automated Framework for Migrating Web applications from SQL to NoSQL database**

Exhibit Session:	Monday B16
Title:	A semi-automated Framework for Migrating Web applications from SQL to NoSQL database
Chair(s):	Rahma Al Mahruqi
Format:	Poster
Description:	We propose a semi-automated approach to migrate highly dynamic SQL-based web applications to ones that uses NoSQL database. There are two parts to this approach: the migration of schema and data, and the migration of the actual application code with embedded queries. Our approach provides contributions to migrating and optimizing the embedded SQL queries to interact with the new database system and changing the application code to use the translated queries.
Industry Company:	IBM
Research Team:	**Queen's University:** Rahma Al Mahruqi Manar Alalfi Thomas Dean

— *Monday B4 (17:45):* **Performance and Security Enhanced Containerized Analytics for IoT Systems**

Exhibit Session:	Monday B4
Title:	Performance and Security Enhanced Containerized Analytics for IoT Systems
Chair(s):	Simon Nadeau , Nathaniel Holeksa , Marios Fokaefs
Format:	Poster
Description:	Our research delves into the pseudo dichotomy that is to be secure or to be performant, in a cloud based IoT architecture. We explore the two in the context of a sensor network in smart buildings, a novel

Give Feedback

	technology that hopes to improve our environmental footprint. In our work, we deploy Spark streaming jobs in Docker containers and explore how to effectively, and automatically, scale them using Docker Compose pursuant to the CPU and memory usage of the standard and the security analytics.
Industry Company:	IBM
Research Team:	**Polytechnique Montreal:** Simon Nadeau Nathaniel Holeksa Marios Fokaefs **York University:** Brian Ramprasad

— *Monday B5 (17:50):* **Randomized Offloading Algorithms for Mobile Cloud Computing**

Exhibit Session:	Monday B5
Title:	Randomized Offloading Algorithms for Mobile Cloud Computing
Chair(s):	Haleh Shahzad , Majid Risman Kar Zadeh
Format:	Poster
Description:	Computational offloading can improve the energy efficiency of mobile devices, by executing some tasks of a mobile application in the cloud. In my research, a set of new algorithms are presented to find the best decision for each task. These algorithms iteratively improve an offloading decision vector, trying to find the optimal one that can minimize the total energy consumption of the mobile device when the application is running while satisfying the delay requirement of the application.
Industry Company:	IBM
Research Team:	**McMaster University:** Haleh Shahzad **Rogers Communications:** Majid Risman Kar Zadeh

— *Monday B6 (17:55):* **Exploring Challenges In Adoption of Continuous Delivery**

Exhibit Session:	Monday B6
Title:	Exploring Challenges In Adoption of Continuous Delivery
Chair(s):	Trevor Rae
Format:	Poster
Description:	We look at the technical and cultural problems that exist in regards to adopting continuous

Give Feedback

	delivery/integration into a company's infrastructure.
Industry Company:	IBM
Product Name:	IBM Cloud Continuous Delivery Service
Research Team:	**University Of Victoria:** Trevor Rae Dr.Damian

— *Monday B7 (18:00):* **Container Technologies! Which One Best Suits Eclipse IoT Deployments?**

Exhibit Session:	Monday B7
Title:	Container Technologies! Which One Best Suits Eclipse IoT Deployments?
Chair(s):	Mohab Aly , Foutse Khomh
Format:	Poster
Description:	The increasing apex in the IoT paradigm requires specific platforms to allow different components to communicate. Our research dives into assessing the performance implications of the Open-Sourced Eclipse IoT, including Hono, to provide guidelines while deploying them in production. In our work, we setup loads in container technologies, such as Kubernetes and OpenShift, and explore how to effectively scale them out using EnMasse while analyzing the CPU usage, memory consumption and network I/O.
Industry Company:	IBM
Research Team:	**Polytechnique Montréal:** Mohab Aly Foutse Khomh

— *Monday B8 (18:05):* **Fall as an irregular pattern in IoT streaming data**

Exhibit Session:	Monday B8
Title:	Fall as an irregular pattern in IoT streaming data
Chair(s):	Sazia Mahfuz
Format:	Poster
Description:	Detecting patterns in real time streaming data has been an interesting and challenging data analytics problem. In this work, we address the problem of learning an irregular human activity pattern, fall, from streaming IoT data using wearable sensors. We present a deep neural network model for detecting fall

Give Feedback

based on accelerometer data using an online dataset "MobiAct" by Vavoulas et al. The initial model was developed using IBM Watson studio and then later transferred and deployed on IBM Cloud.

Industry Company: IBM

Research Team: **Queen's University:** Sazia Mahfuz Dr. Farhana Zulkernine
IBM Canada: Peter Nicholls

Short Talk Sessions start from 18:10 (10 mins each). Talks order as follows:

— *S1 (18:10):* **Db2 on Cloud Free Forever Cloud Database**

Booth Number: S1

Title: Db2 on Cloud Free Forever Cloud Database

Chair(s): Saygin Yag , Simon Lightstone

Format: Short Talk

Description: In 2018, Db2 on Cloud introduced a lite plan that is completely free for as long as you need it. This plan provides users with 100MB of cloud database space on a multi-tenant system. This presents a perfect starting point for students, university classes, idea-stage startups and anyone who wants to have a world-class cloud database, for free!

Business unit IBM Analytics

Product Name: Db2 on Cloud

Web Site: https://www.ibm.com/cloud/db2-on-cloud

Research Team: **IBM:** Saygin Yag - Offering Manager, Db2 on Cloud Simon Lightstone - Offering Manager, Db2 on

— *S2 (18:20):* **User Research: Mapping the end-to-end experience of a System Administrator’s deployment**

Booth Number: S2

Title: User Research: Mapping the end-to-end experience of a System Administrator’s deployment

Chair(s): Arun Martin , Steve Murphy

Format: Short Talk

Give Feedback

Description:	In this talk, we share how design research helped the product stakeholders identify issues in the end-to-end experiences and various support channels accessed while deploying complex enterprise applications. During this design research journey, the product stakeholders understood the ecosystem of primary & secondary users, dependencies and handoffs and how the current customer journey impacted the organizational teams and processes. The findings from the study informed the creation of a MVP.
Research Team:	**IBM Canada Lab:** Steve Murphy, UX Design Lead Arun Joseph, Design Researcher

— *S3 (18:30):* Assessing Databases for GDPR Compliance

Booth Number:	S3
Title:	Assessing Databases for GDPR Compliance
Chair(s):	Devan Shah , Walid Rjaibi , Larry Lindsay , Josue Diaz
Format:	Short Talk
Description:	The General Data Protection Regulation (GDPR) requires organization to implement adequate controls to protect Personally Identifiable Information (PII). A critical step in that journey is impact assessment. That is, understanding where data subject to GDPR is located and how vulnerable it is. In this session, we will introduce IBM Security Guardium Analyzer, a new SaaS offering which intends to help organizations easily complete the impact assessment for their databases.
Industry Company:	IBM
Business unit	IBM Security
Product Name:	IBM Security Guardium Analyzer
Web Site:	https://www.ibm.com/marketplace/guardium-analyzer
Research Team:	**IBM:** Devan Shah Walid Rjaibi Larry Lindsay Josue Diaz

— *S4 (18:40):* A Dynamic Knowledge Management System for Decision Support

Booth Number:	S4
Title:	A Dynamic Knowledge Management System for Decision Support
Chair(s):	Ftoon Kedwan , Farhana Zulkernine

Give Feedback

Format:	Short Talk
Description:	We focus on; a) studying the state-of the art KMS to select the right system to store hybrid medical knowledge, b) define a multi-level graph structure with schema and metadata to dynamically and efficiently link knowledge sources and also to contain frequently accessed knowledge, and c) develop machine learning models to learn query patterns for efficient knowledge organization and retrieval. We are still at the very early phase of the research. Fig. 1 shows the conceptual architecture.
Product Name:	Knowledge Management System
Web Site:	https://github.com/isahqueens/multi-level_streaming_analytics/blob/master/KMS.jpeg
Research Team:	**Queen's University:** Ftoon Kedwan

— *S5 (18:50):* Distribution and Severity of Test Smells: An exploratory study

Booth Number:	S5
Title:	Distribution and Severity of Test Smells: An exploratory study
Chair(s):	Mohamed Wiem Mkaouer
Format:	Short Talk
Description:	Similar to production code, unit tests are also susceptible to bad programming practices, which can, in turn, have a negative impact not only on the quality of the software system but also on maintenance activities. In this talk, I showcase various test smell types, demonstrate their existence in Java projects, analyze their impact on software maintenance in terms of code vulnerability to changes and bugs.
Industry Company:	IBM
Web Site:	https://testsmells.github.io/
Research Team:	**Rochester Institute of Technology:** Mohamed Wiem

Give Feedback

Expo

Think: Cognitive Computing, Big Data, Cloud, Security and Privacy

October 29 - 31

CASCON News

Download the CASCON 2018 Program Guide

See the CASCON 2018 Proceeding in ACM

Get certain Workshops materials

Call for Papers for CASCON 2019

Give Feedback

The Technology Expo at CASCON 2018 will provide an excellent opportunity to highlight your research results and generate interest in upcoming product and product areas. Exhibits on emerging technologies are encouraged. The format of an exhibit may be a demonstration, poster presentation, or a short talk. A booth is provided at no cost to the exhibitors. Each demonstration booth will be

set up with power and wireless internet access. Exhibitors are responsible for bringing the necessary equipment such as machines, laptops, monitors, power cables, and security locks; and any printed materials including posters.

Accepted exhibits

Congratulations to the exhibit chairs and thank you for your submissions. See you all at CASCON 2018.

Monday | **Tuesday** | Wednesday

Exhibit Sessions - 12:00 to 13:30 and 17:30 to 19:00

— *B1:* Foodie: Conversational AI for the smart kitchen

Booth Number:	B1
Title:	Foodie: Conversational AI for the smart kitchen
Chair(s):	Prashanti Priya Angara , Miguel Jimenez , Hausi Muller , Ulrike Stege
Format:	Demo
Description:	Foodie is a cognitive text-and-voice based conversational agent that augments the capabilities of home cooks by incorporating health-related information to aid one's eating habits. This exhibit demonstrates the capabilities and limitations of existing conversational frameworks via a hands-on demo of Foodie which is built with IBM's Natural Language Understanding and voice services. On the back-end, Foodie is connected to food databases such as Spoonacular (for recipes).
Industry Company:	IBM
Research Team:	**University of Victoria:** Prashanti Priya Angara, Hausi A. Müller, Ulrike Stege

— *B2:* Privacy-aware Efficient Visual Recognition Services for Smart Laboratories

Booth Number:	B2

Give Feedback

Title:	Privacy-aware Efficient Visual Recognition Services for Smart Laboratories
Chair(s):	Christianne Huber , Mahima Chaudhary , Joydeep Mukherjee , Marin Litoiu
Format:	Demo
Description:	With concerns about personal privacy on the rise, individuals may prefer that their images captured by visual recognition services have controlled privacy levels. IBM Watson Visual Recognition Service ensures efficiency but individuals may prefer not to send images to such a public service. We propose an intelligent privacy-aware framework that maintains efficiency while delivering a desired privacy level to its end users by dynamically directing images to a localized privacy guaranteed server.
Industry Company:	IBM
Product Name:	IBM Watson Visual Recognition Service
Research Team:	**York University:** Christianne Huber Mark Shtern Joydeep Mukherjee Marin Litoiu **Dayalbagh Educational Institute:** Mahima Chaudhary **IBM:** Vio Onut

— *B3:* Data Science 2 - Bootcamp Materials

Booth Number:	B3
Title:	Data Science 2 - Bootcamp Materials
Chair(s):	Marcellus Mindel , Dennis Buttera , Jennifer Collins
Format:	Poster
Description:	In the second of three exhibits on Data Science Education, we will discuss uses for 12 hours of course materials that cover the data science workflow using IBM Watson Studio. These materials are suitable for inclusion as part of an introductory Data Science course. They could also be delivered as an intensive hands-on bootcamp, for example, to complement a business program, or to enable graduate students from different backgrounds to begin with a consistent practical foundation for data science.
Industry Company:	IBM
Research Team:	**IBM Advanced Studies:** Marcellus Mindel Dennis Buttera Jennifer Collins

Give Feedback

— *B4:* In-Depth Evaluation of Task Parallelism in Node.js

Booth Number:	B4
Title:	In-Depth Evaluation of Task Parallelism in Node.js
Chair(s):	Maria Patrou
Format:	Poster
Description:	Node.js is a framework for asynchronous I/O, event-driven, server-side Javascript. Computationally intensive tasks are bound to the performance of a single core. We evaluate the performance of some representative Node.js multi-processed and multi-threaded techniques used for scaling, running a CPU-intensive task. Computation, memory and garbage collection metrics, the effect of the execution environment and similarities between these modules are investigated and recommendations are provided.
Industry Company:	IBM
Project Number	1002
Product Name:	IBM SDK for Node.js
Web Site:	http://www.unb.ca/research/casatlantic/index.html
Research Team:	**University of New Brunswick:** Maria Patrou Kenneth B. Kent **IBM Canada:** Michael Dawson

— *B5:* Supporting Software Variability in Eclipse OMR

Booth Number:	B5
Title:	Supporting Software Variability in Eclipse OMR
Chair(s):	Samer AL Masri
Format:	Poster
Description:	Eclipse OMR is a C++ library of language runt-ime components that supports multiple architectures and various languages. Aiming to support variability in OMR, we created OMRStatistics, a static analysis tool that helps developers collect statistics about the project and make development decisions about

Give Feedback

variability. I will be explaining how we use OMRStatistics to support variability in OMR and our experience switching from static polymorphism to dynamic polymorphism in OMR's implementation.

Industry Company: IBM

Web Site: https://github.com/samasri/omr/tree/master/tools/compiler/OMRStatistics

Research Team: **University of Alberta:** Samer

— *B6:* Fear and Disgust during the 2015 Canadian Election: Automated Analysis with Watson

Booth Number: B6

Title: Fear and Disgust during the 2015 Canadian Election: Automated Analysis with Watson

Chair(s): Trevor Deley

Format: Poster

Description: There is a tremendous opportunity for automated content analysis solutions in communications research as well as the media industry at large. This poster presents a proof of concept for automated content analysis in a media studies context. Currently there is a gap between the long standing field of content analysis in media studies and the field of sentiment analysis that arose as a subset of natural language processing. This exhibit hopes to begin to bring the two fields closer together.

Industry Company: IBM

Research Team: **University of Ottawa:** Trevor Deley

— *B7:* ROAD - Realtime Outlier and Anomaly Detection for IBM Cloud

Booth Number: B7

Title: ROAD - Realtime Outlier and Anomaly Detection for IBM Cloud

Chair(s): William Pourmajidi , Lei Zhang , Andriy Miranskyy

Format: Poster

Description: We refer to logs when there are errors or anomalies. Hence, anomaly detection is a crucial part of the log analysis, and if anomalies are detected at an early stage, many issues can be prevented. In this poster, we

Give Feedback

	discuss practical issues and potential solutions needed to create an anomaly detection solution tailored for a large-scale Cloud platform consisting of multiple software and hardware components.
Industry Company:	Ryerson University
Product Name:	IBM Cloud
Research Team:	**Ryerson University:** William Pourmajidi, Lei Zhang, Andriy Miran **IBM:** Tony Erwin, John Steinbacher, David Godwin, David Kalmuk, Aslam Nomani, Vivien Page, Nic

— *B8:* Reconstruct the genealogy of phishing kits source code modifications

Booth Number:	B8
Title:	Reconstruct the genealogy of phishing kits source code modifications
Chair(s):	Stéphane Heudron
Format:	Poster
Description:	In this project, we will study the similarity between phishing kits belonging to a reference database and develop an approach to extract a plausible genealogy to explain source code modifications. Our initial approach is based on static analysis of PHP scripts. To investigate the similarity, we will explore several distances and algorithms responding to our constraints (incrementality, efficiency). We will present preliminary results of similarity between kits and some candidate genealogy trees.
Industry Company:	IBM
Project Number	1059
Product Name:	IBM Forensics
Research Team:	**Polytechnique Montréal:** Stephane Heudron Ettore Merlo **University of Ottawa:** Guy-Vincent Jourdan **IBM Canada:** Vio Onut

— *B9:* Machine Learning to Provide Intelligent Digital Health for Underserved Regions

Booth Number:	B9
Title:	Machine Learning to Provide Intelligent Digital Health for Underserved Regions

Give Feedback

Chair(s):	Kamran Sartipi , Jing Chong
Format:	Poster
Description:	This project provides advanced information intelligence for physicians in underserved regions to utilize mined knowledge of medical specialties. A pilot project is ongoing with collaboration of Gastroenterologists and Family Physicians at McMaster University. Resources are: anonymized specialty patient data; research datasets; expert decisions of specialists. Concept lattice, machine learning will produce RDF knolwlegebase. Intelligent consultant service navigates RDF and interacts with doctors.
Industry Company:	IBM
Web Site:	http://profs.degroote.mcmaster.ca/ads/sartipi/HTMLfiles/Research/ResearchProj2.php
Research Team:	**McMaster and East Carolina Universities:** Kamran Sartipi (CS at ECU and IS at McMaster David Armstrong (Gastroenterology at McMaste Smita Halder (Gastroenterology at McMaster U Henry Siu (Family Medicine at McMaster Unive Jing Chong (Information Systems at McMaster

— *B10:* Detecting Misuses of Crypto APIs

Booth Number:	B10
Title:	Detecting Misuses of Crypto APIs
Chair(s):	Kristen Newbury
Format:	Poster
Description:	This exhibit will present the integration of cryptography API misuse detection into a Just-in-Time context. My goal for a novel contribution to this research will be to enable a JIT compiler to perform security checks that have been generated by a pre-existing API misuse detection framework. Such an integration is expected to benefit the precision and performance of the security analyses, since runtime values can be exploited in this context.
Industry Company:	IBM
Project Number	1064
Product Name:	IBM SDK for Java, Testarossa, IBM JIT Compiler
Research Team:	**University of Alberta:** Kristen

Give Feedback

— *B11:* Pressure Sensitive Mats for Patient Monitoring in the NICU

Booth Number:	B11
Title:	Pressure Sensitive Mats for Patient Monitoring in the NICU
Chair(s):	Shermeen Nizami , James R. Green
Format:	Demo
Description:	A research project that is evaluating the use of the pressure sensitive mat (PSM) technology as well RGB-D streaming video as continuous patient monitoring modalities. This research is conducted at the Neonatal Intensive Care Unit (NICU) at CHEO, Ottawa. We have developed and deployed two software apps at the patient bedside to collect physiologic and clinical event data from patients in real-time. We have acquired data from fifteen critically ill newborn babies up to date.
Product Name:	Watson Analytics and Health
Research Team:	**Carleton University:** Shermeen Nizami James R. Green Randy Giffen

— *B12:* Susereum - A Blockchain Platform for Sustainable Software Systems

Booth Number:	B12
Title:	Susereum - A Blockchain Platform for Sustainable Software Systems
Chair(s):	Omar Badreddin
Format:	Demo
Description:	Software sustainability is a systematic challenge, and not an individual or team failings. This systematic challenge stems because engineers often take shortcuts to maximize immediate gains that become detrimental to software longevity. We demo Susereum, a blockchain platform that creates a balance between long term longevity and short term necessities. Susereum creates a permanent public ledger to monitory the sustainability of software code modifications and credit their authors.
Industry Company:	Susereum
Product Name:	Susereum
Web Site:	https://susereum.com/
Research Team:	**University of Texas:** Omar Badreddin

Give Feedback

— *B13:* Localizing Error Prone Code Using Software Analytics

Booth Number:	B13
Title:	Localizing Error Prone Code Using Software Analytics
Chair(s):	KONSTANTINOS KONTOGIANNIS , Sanjay Ghanathey , Marios Grigoriou
Format:	Demo
Description:	- Evaluation study into how software analytics and machine learning can be used to predict whether a source code file contains one or more errors that may contribute to a major system failure. - The approach is utilizing information extracted both from the system's source code, such as code metrics, and from a series of DevOps tools. - Results indicate that data obtained from repositories and used as training features performed equal to or better than source code metrics
Industry Company:	IBM
Project Number	1048
Product Name:	DevOps Services, DevOps Insights
Research Team:	**Western University:** KONSTANTINOS KONTOGIANNIS SANJAY GHANATHEY MARIOS GRIGORIOU **IBM US:** ALBERTO GIAMMARIA BRIAN COUNTRYMAN **IBM CANADA:** CHRIS BREALEY

— *B14:* Eclipse OMR on the Embedded AArch64 Platform

Booth Number:	B14
Title:	Eclipse OMR on the Embedded AArch64 Platform
Chair(s):	Jean-Philippe Legault , Aaron Graham
Format:	Poster
Description:	The AArch64 architecture—a subset of the ARM architectural model—has the characteristics to be driving the next generation of data centers. With low power and cost overhead, AArch64 helps move computation closer to the originating data; following the “fog computing” model. We develop support for AArch64 in the Eclipse OMR Toolkit to evaluate modern runtimes with new computing paradigms.

Give Feedback

Industry Company:	IBM
Project Number	1035
Product Name:	Eclipse OMR
Web Site:	http://www.unb.ca/research/casatlantic/index.html
Research Team:	**University of New Brunswick:** Jean-Philippe Legault Aaron G. Graham Kenneth B. Kent **IBM:** Daryl Maier James Kingdon

— *B15:* (Semi)Automatic Construction of Access-Controlled Web Data Services

Booth Number:	B15
Title:	(Semi)Automatic Construction of Access-Controlled Web Data Services
Chair(s):	Kalvin Eng
Format:	Demo
Description:	A demo of a prototype system implementing a novel fitness application use case. This system uses model-driven methodology to (1) automatically abstract the representation of relational database schemas into a visual model and map them to existing ontologies, (2) use the mappings in order to create different role-based access-control views of APIs, and (3) automatically generate API endpoints and their responses, based on these mappings.
Industry Company:	IBM
Research Team:	**University of Alberta:** Kalvin Eng Diego Serrano Eleni Stroulia Jacob Jaremko

— *B16:* Design Patterns for Machine Learning Solutions

Booth Number:	B16
Title:	Design Patterns for Machine Learning Solutions
Chair(s):	Soroosh Nalchigar
Format:	Demo
Description:	dvanced analytics and machine learning algorithms have a wide and rapidly growing range of applications in business contexts. In this project, we apply a design patterns approach to reduce the level of expertise

Give Feedback

	and efforts needed to develop machine learning solutions to be incorporated into cognitive business operations. The approach is based on design knowledge catalogs encoded according to a meta-model. Implementation of the approach is being investigated in the context of a BPM suite.
Industry Company:	IBM
Project Number	1030
Product Name:	Business Process Management
Research Team:	**University of Toronto:** Soroosh Nalchigar Eric Yu **IBM Canada:** Sebastian Carbajales Yazan Obeidi John Green Allen Chan

— *B17:* Open Liberty and Eclipse MicroProfile: Foundation for Microservices

Booth Number:	B17
Title:	Open Liberty and Eclipse MicroProfile: Foundation for Microservices
Chair(s):	Eveline Cai , Gilbert Kwan , YK Chang , Panagiotis Roubatsis
Format:	Poster
Description:	We will introduce Open Liberty, which is the largest open source contributions ever done by IBM, what it is, and why Open Liberty. We invite you to come see what Open Liberty is, collaborate with us, and contribute to this fast, dynamic, and simply amazing application server! Eclipse MicroProfile is the next step in the evolution of enterprise Java technologies for the microservices and cloud native world.
Business unit	IBM Hybrid Cloud
Product Name:	Open Liberty
Research Team:	**IBM:** Eveline Cai Gilbert Kwan Yee-Kang Chang Panagiotis Roubatsis

— *B18:* Node.js Scalability Investigation in the Cloud

Booth Number:	B18
Title:	Node.js Scalability Investigation in the Cloud
Chair(s):	Jiapeng Zhu , Maria Patrou
Format:	Poster

Give Feedback

Description:	Node.js has gained popularity in cloud development. However, scalability issues can limit the number of concurrent requests while achieving an acceptable level of performance. We design and implement Ibenchjs, a scalability-oriented benchmarking framework, and a set of sample test applications. We deploy Ibenchjs in a local and isolated cloud to report scalability-related measurements and issues of Node.js as well as performance bottlenecks to investigate Node.js scalability issues in the cloud.
Industry Company:	IBM
Project Number	1002 Node.js
Product Name:	IBM SDK for Node.js
Web Site:	http://www.unb.ca/research/casatlantic/index.html
Research Team:	**University of New Brunswick:** Jiapeng Zhu Panagiotis Patros Kenneth B. Kent **IBM Runtime Technologies:** Michael Dawson

— *B19:* Feasibility of Internal Object Pools to Reduce Memory Management Activity

Booth Number:	B19
Title:	Feasibility of Internal Object Pools to Reduce Memory Management Activity
Chair(s):	Konstantin Nasartschuk
Format:	Poster
Description:	Object pools is a software engineering pattern used to reuse object instances without the need of repeated allocation and instantiation. While the benefits of using object pool structures are still present when used in a garbage collected environment, it adds a memory management component to the development process. This work investigates automatically created object pools for predefined classes. Automatic object pools are implemented and discussed using the GenCon GC and Balanced GC policies.
Industry Company:	IBM
Project Number	1019
Product Name:	Object Locality (IBM SDK)
Research Team:	**University of New Brunswick:** Konstantin Nasartschuk Kenneth B. Kent Stephen A. MacKay **IBM Canada:** Aleksandar Micic

Give Feedback

— *B20:* Cold Object Identification, Sequestrating and Revitalization

Booth Number:	B20
Title:	Cold Object Identification, Sequestrating and Revitalization
Chair(s):	Abhijit Taware
Format:	Poster
Description:	Cold objects are alive and infrequently accessed. Such objects are an overhead for garbage collection. Identifying and moving them to a separate memory space would benefit the real memory footprint and cache coherence. Intel's NVRam technology will be used as a cold heap for this research. Application will be profiled to gather statistics about object temperature. This will be utilized for cold object sequestration and revitalization.
Industry Company:	IBM
Project Number	1001
Product Name:	Eclipse OMR, IBM Java SDK, OMR, J9
Web Site:	http://www.unb.ca/research/casatlantic/projects/datatemp.html
Research Team:	**UNB:** Abhijit Taware Kenneth B. Kent Gerhard W. Dueck **IBM:** Charlie Gracie

— *B21:* Persistent Memory Storage of Cold Regions in the OpenJ9 Java Virtual Machine

Booth Number:	B21
Title:	Persistent Memory Storage of Cold Regions in the OpenJ9 Java Virtual Machine
Chair(s):	Scott Young
Format:	Poster
Description:	In this poster an optimization technique for language runtimes with automatic memory management is presented. The technique involves segregating objects into different memory areas, backed by different memory devices, on a per object basis. This technique is compared against operating system paging mechanisms with swap partitions. Two different schemes for determining which objects should be segregated into slower memory are tested and their results are discussed.

Give Feedback

Industry Company:	IBM
Project Number	1001
Product Name:	Eclipse OMR, IBM Java SDK, OMR, J9
Web Site:	http://www.unb.ca/research/casatlantic/projects/datatemp.html
Research Team:	**Centre for Advanced Studies - Atlantic:** Scott Young Kenneth B. Kent Gerhard Dueck **IBM:** Charlie Gracie

— *B22:* A novel approach to generational garbase collection

Booth Number:	B22
Title:	A novel approach to generational garbase collection
Chair(s):	Kim Briggs
Format:	Poster
Description:	Presenting a new algorithm for evacuating live objects from the nursery during generational collection cycles. This algorithm uses a stack to follow object reference chains and selectively copies small objects close to referring pointers, thereby improving collector throughput and object locality within the generational heap. This algorithm is being developed within IBMs open source OMR framework, which provides language-agnostic components to provide common components for language runtimes.
Industry Company:	IBM
Web Site:	https://github.com/eclipse/omr
Research Team:	**IBM Canada:** Kim Briggs

— *B23:* Math function performance acceleration via integrated HW and algorithm design

Booth Number:	B23
Title:	Math function performance acceleration via integrated HW and algorithm design
Chair(s):	Christopher Anand , Lucas Dutton , Adele Olejarz , Robert Enenkel
Format:	Poster
Description:	We present results from our CAS project to design HW instructions together with algorithms to work

Give Feedback

	together to produce fast, accurate mathematical functions.
Industry Company:	IBM
Project Number	1006
Product Name:	POWER processors, MASS math libraries
Web Site:	https://www.eng.mcmaster.ca/cas/people/faculty/christopher-anand
Research Team:	**IBM:** Robert Enenkel **McMaster University:** Christopher Anand Lucas Dutton Adele Olejarz

— ***B24:*** **Automating the Detection of Third-Party Java Library Migration At The Function Level**

Booth Number:	B24
Title:	Automating the Detection of Third-Party Java Library Migration At The Function Level
Chair(s):	Mohamed Wiem Mkaouer
Format:	Poster
Description:	This poster complements our accepted paper entitled "Automating the Detection of Third-Party Java Library Migration At The Function Level". In this poster, we show the design of algorithms related to mining all the manually-performed function replacements for a given library migration. We further show the detailed results of the comparative study between state-of-art approaches for detecting migration traces at the function level.
Industry Company:	IBM
Web Site:	http://migrationlab.net/
Research Team:	**Rochester Institute of Technology:** Mohamed Wiem

Expo Reception / Short Talks / Expo Presentations

Short Talk Sessions starts from 17:30 (5 mins each). Presentations order as follows:

— *Tuesday B1 (17:30):* **Foodie: Conversational AI for the smart kitchen**

Give Feedback

Exhibit Session:	Tuesday B1
Title:	Foodie: Conversational AI for the smart kitchen
Chair(s):	Prashanti Priya Angara , Miguel Jimenez , Hausi Muller , Ulrike Stege
Format:	Demo
Description:	Foodie is a cognitive text-and-voice based conversational agent that augments the capabilities of home cooks by incorporating health-related information to aid one's eating habits. This exhibit demonstrates the capabilities and limitations of existing conversational frameworks via a hands-on demo of Foodie which is built with IBM's Natural Language Understanding and voice services. On the back-end, Foodie is connected to food databases such as Spoonacular (for recipes).
Industry Company:	IBM
Research Team:	**University of Victoria:** Prashanti Priya Angara, Hausi A. Müller, Ulrike Stege

— *Tuesday B2 (17:35):* **Privacy-aware Efficient Visual Recognition Services for Smart Laboratories**

Exhibit Session:	Tuesday B2
Title:	Privacy-aware Efficient Visual Recognition Services for Smart Laboratories
Chair(s):	Christianne Huber , Mahima Chaudhary , Joydeep Mukherjee , Marin Litoiu
Format:	Demo
Description:	With concerns about personal privacy on the rise, individuals may prefer that their images captured by visual recognition services have controlled privacy levels. IBM Watson Visual Recognition Service ensures efficiency but individuals may prefer not to send images to such a public service. We propose an intelligent privacy-aware framework that maintains efficiency while delivering a desired privacy level to its end users by dynamically directing images to a localized privacy guaranteed server.
Industry Company:	IBM
Product Name:	IBM Watson Visual Recognition Service
Research Team:	**York University:** Christianne Huber Mark Shtern Joydeep Mukherjee Marin Litoiu **Dayalbagh Educational Institute:** Mahima Chaudhary **IBM:** Vio Onut

Give Feedback

— *Tuesday B16 (17:40):* Design Patterns for Machine Learning Solutions

Exhibit Session:	Tuesday B16
Title:	Design Patterns for Machine Learning Solutions
Chair(s):	Soroosh Nalchigar
Format:	Demo
Description:	dvanced analytics and machine learning algorithms have a wide and rapidly growing range of applications in business contexts. In this project, we apply a design patterns approach to reduce the level of expertise and efforts needed to develop machine learning solutions to be incorporated into cognitive business operations. The approach is based on design knowledge catalogs encoded according to a meta-model. Implementation of the approach is being investigated in the context of a BPM suite.
Industry Company:	IBM
Project Number	1030
Product Name:	Business Process Management
Research Team:	**University of Toronto:** Soroosh Nalchigar Eric Yu **IBM Canada:** Sebastian Carbajales Yazan Obeidi John Green Allen Chan

— *Tuesday B4 (17:45):* In-Depth Evaluation of Task Parallelism in Node.js

Exhibit Session:	Tuesday B4
Title:	In-Depth Evaluation of Task Parallelism in Node.js
Chair(s):	Maria Patrou
Format:	Poster
Description:	Node.js is a framework for asynchronous I/O, event-driven, server-side Javascript. Computationally intensive tasks are bound to the performance of a single core. We evaluate the performance of some representative Node.js multi-processed and multi-threaded techniques used for scaling, running a CPU-intensive task. Computation, memory and garbage collection metrics, the effect of the execution environment and similarities between these modules are investigated and recommendations are provided.

Give Feedback

Industry Company:	IBM
Project Number	1002
Product Name:	IBM SDK for Node.js
Web Site:	http://www.unb.ca/research/casatlantic/index.html
Research Team:	**University of New Brunswick:** Maria Patrou Kenneth B. Kent **IBM Canada:** Michael Dawson

— *Tuesday B5 (17:50):* **Supporting Software Variability in Eclipse OMR**

Exhibit Session:	Tuesday B5
Title:	Supporting Software Variability in Eclipse OMR
Chair(s):	Samer AL Masri
Format:	Poster
Description:	Eclipse OMR is a C++ library of language runt-ime components that supports multiple architectures and various languages. Aiming to support variability in OMR, we created OMRStatistics, a static analysis tool that helps developers collect statistics about the project and make development decisions about variability. I will be explaining how we use OMRStatistics to support variability in OMR and our experience switching from static polymorphism to dynamic polymorphism in OMR's implementation.
Industry Company:	IBM
Web Site:	https://github.com/samasri/omr/tree/master/tools/compiler/OMRStatistics
Research Team:	**University of Alberta:** Samer

— *Tuesday B6 (17:55):* **Fear and Disgust during the 2015 Canadian Election: Automated Analysis with Watson**

Exhibit Session:	Tuesday B6
Title:	Fear and Disgust during the 2015 Canadian Election: Automated Analysis with Watson
Chair(s):	Trevor Deley
Format:	Poster
Description:	There is a tremendous opportunity for automated content analysis solutions in communications research

Give Feedback

as well as the media industry at large. This poster presents a proof of concept for automated content analysis in a media studies context. Currently there is a gap between the long standing field of content analysis in media studies and the field of sentiment analysis that arose as a subset of natural language processing. This exhibit hopes to begin to bring the two fields closer together.

Industry Company: IBM

Research Team: **University of Ottawa:** Trevor Deley

— *Tuesday B7 (18:00):* ROAD - Realtime Outlier and Anomaly Detection for IBM Cloud

Exhibit Session: Tuesday B7

Title: ROAD - Realtime Outlier and Anomaly Detection for IBM Cloud

Chair(s): William Pourmajidi , Lei Zhang , Andriy Miranskyy

Format: Poster

Description: We refer to logs when there are errors or anomalies. Hence, anomaly detection is a crucial part of the log analysis, and if anomalies are detected at an early stage, many issues can be prevented. In this poster, we discuss practical issues and potential solutions needed to create an anomaly detection solution tailored for a large-scale Cloud platform consisting of multiple software and hardware components.

Industry Company: Ryerson University

Product Name: IBM Cloud

Research Team: **Ryerson University:** William Pourmajidi, Lei Zhang, Andriy Miran

IBM: Tony Erwin, John Steinbacher, David Godwin, David Kalmuk, Aslam Nomani, Vivien Page, Nic

— *Tuesday B8 (18:05):* Reconstruct the genealogy of phishing kits source code modifications

Exhibit Session: Tuesday B8

Title: Reconstruct the genealogy of phishing kits source code modifications

Chair(s): Stéphane Heudron

Format: Poster

Description: In this project, we will study the similarity between phishing kits belonging to a reference database and

Give Feedback

develop an approach to extract a plausible genealogy to explain source code modifications. Our initial approach is based on static analysis of PHP scripts. To investigate the similarity, we will explore several distances and algorithms responding to our constraints (incrementality, efficiency). We will present preliminary results of similarity between kits and some candidate genealogy trees.

Industry Company: IBM
Project Number 1059
Product Name: IBM Forensics
Research Team: **Polytechnique Montréal:** Stephane Heudron Ettore Merlo
University of Ottawa: Guy-Vincent Jourdan
IBM Canada: Vio Onut

— *Tuesday B9 (18:10):* **Machine Learning to Provide Intelligent Digital Health for Underserved Regions**

Exhibit Session: Tuesday B9
Title: Machine Learning to Provide Intelligent Digital Health for Underserved Regions
Chair(s): Kamran Sartipi , Jing Chong
Format: Poster
Description: This project provides advanced information intelligence for physicians in underserved regions to utilize mined knowledge of medical specialties. A pilot project is ongoing with collaboration of Gastroenterologists and Family Physicians at McMaster University. Resources are: anonymized specialty patient data; research datasets; expert decisions of specialists. Concept lattice, machine learning will produce RDF knolwlegebase. Intelligent consultant service navigates RDF and interacts with doctors.
Industry Company: IBM
Web Site: http://profs.degroote.mcmaster.ca/ads/sartipi/HTMLfiles/Research/ResearchProj2.php
Research Team: **McMaster and East Carolina Universities:** Kamran Sartipi (CS at ECU and IS at McMaster David Armstrong (Gastroenterology at McMaste Smita Halder (Gastroenterology at McMaster U Henry Siu (Family Medicine at McMaster Unive Jing Chong (Information Systems at McMaster

Give Feedback

— *Tuesday B10 (18:15):* **Detecting Misuses of Crypto APIs**

Exhibit Session:	Tuesday B10
Title:	Detecting Misuses of Crypto APIs
Chair(s):	Kristen Newbury
Format:	Poster
Description:	This exhibit will present the integration of cryptography API misuse detection into a Just-in-Time context. My goal for a novel contribution to this research will be to enable a JIT compiler to perform security checks that have been generated by a pre-existing API misuse detection framework. Such an integration is expected to benefit the precision and performance of the security analyses, since runtime values can be exploited in this context.
Industry Company:	IBM
Project Number	1064
Product Name:	IBM SDK for Java, Testarossa, IBM JIT Compiler
Research Team:	**University of Alberta:** Kristen

Short Talk Sessions start from 18:20 (10 mins each). Talks order as follows:

— *S1 (18:20):* **Discovering business rules automatically from legacy applications**

Booth Number:	S1
Title:	Discovering business rules automatically from legacy applications
Chair(s):	Arun Martin , David Budreau , Tala El Hallak
Format:	Short Talk
Description:	The BRM research and development team will demonstrate how our revised architecture focusing on machine learning and user research has transformed last year's BRM preview in Application Discovery and Delivery Intelligence (ADDI) to enable automated discovery of terms and to attach them to relevant code snippets through a collaborative workflow between business and information technology teams resulting in increased knowledge, reach, and ability for members to make faster business decisions.

Give Feedback

Product Name:	Business Rules Discovery (ADDI)
Research Team:	**IBM:** Arun Joseph David Budreau Tala Hallak

— *S2 (18:30):* REST APIs with LoopBack 4 and OpenAPI 3

Booth Number:	S2
Title:	REST APIs with LoopBack 4 and OpenAPI 3
Chair(s):	Taranveer Virk , Diana Lau , Biniam Admikew , Janny Hou
Format:	Short Talk
Description:	REST APIs power web applications, mobile applications, IoT devices & more. LoopBack 4 is an API creation framework that has been rewritten from the ground up and features a new programming interface, OpenAPI support, Dependency Injection, a new CLI and so much more. LoopBack makes it easy to create powerful OpenAPI 3.x (an industry standard) compliant REST APIs in minutes. It supports connecting to multiple SQL and NoSQL databases. Poster will show how to get started in minutes with LoopBack 4.
Business unit	Hybrid Cloud
Product Name:	LoopBack
Web Site:	http://v4.loopback.io/
Research Team:	**IBM:** Taranveer Virk Diana Lau Biniam Admikew Janny Hou

— *S3 (18:40):* Dynamic Large-Scale Graph Processing over Data Streams

Booth Number:	S3
Title:	Dynamic Large-Scale Graph Processing over Data Streams
Chair(s):	Tariq Abughofa
Format:	Short Talk
Description:	Processing time-evolving graphs is still a challenging problem. In this work, we studied approaches to dynamic graph processing with Spark. The study helped us design and develop Sprouter, a dynamic graph processing framework that enables storing enormous graphs, real-time updates, and efficient OLTP and

Give Feedback

analytics. We selected community detection as a case study of incremental analytics with Sprouter and proposed IDWCC, a novel community detection algorithm for dynamic graphs.

Industry Company: IBM

Research Team: **Queen's University:** Tariq

— *S4 (18:50):* A Multilevel Streaming Data Analytics Infrastructure for Predictive Analytics

Booth Number: S4

Title: A Multilevel Streaming Data Analytics Infrastructure for Predictive Analytics

Chair(s): Haruna Isah , Farhana Zulkernine

Format: Short Talk

Description: We are developing a real-time multilevel streaming data analytics infrastructure using cutting-edge streaming data processing engines, in-memory data structure and analytics tools which will preprocess, filter and store necessary data in the memory for higher level processing. This talk is aimed at showcasing our preliminary framework and is targeted at individuals and organizations that are planning to optimize their current infrastructure to embrace multilevel streaming data analytics.

Industry Company: Gnowit Inc

Product Name: Multilevel Streaming Infrastructure for News Analytics

Web Site: https://github.com/isahqueens/multi-level_streaming_analytics

Research Team: **Queens University:** Haruna Isah Farhana Zulkernine

Give Feedback

Expo

Think: Cognitive Computing, Big Data, Cloud, Security and Privacy

October 29 - 31

CASCON News

Download the CASCON 2018 Program Guide

See the CASCON 2018 Proceeding in ACM

Get certain Workshops materials

Call for Papers for CASCON 2019

Give Feedback

The Technology Expo at CASCON 2018 will provide an excellent opportunity to highlight your research results and generate interest in upcoming product and product areas. Exhibits on emerging technologies are encouraged. The format of an exhibit may be a demonstration, poster presentation, or a short talk. A booth is provided at no cost to the exhibitors. Each demonstration booth will be set up with power and wireless internet access. Exhibitors are responsible for bringing the necessary equipment such as machines, laptops, monitors, power cables, and security locks; and any printed materials including posters.

Accepted exhibits

Congratulations to the exhibit chairs and thank you for your submissions. See you all at CASCON 2018.

Monday	Tuesday	**Wednesday**

Exhibit Sessions - 08:30 to 10:45 and 12:00 to 13:30

— *B1:* Model based Interference-aware Management System for Microservice Architectures

Booth Number:	B1
Title:	Model based Interference-aware Management System for Microservice Architectures
Chair(s):	Joydeep Mukherjee , Marin Litoiu , Cornel Barna
Format:	Demo
Description:	The performance of microservices deployed in containers on a cloud platform can often be impacted due to interference. Interference is caused by other applications that run concurrently with the microservice containers on the same host and compete for host level shared resources. We build a model based management system that takes into account the effect of interference and works in parallel with existing platforms such as IBM Cloud Automation Manager for efficient deployment of microservices.
Industry Company:	IBM
Product Name:	Cloud Automation Manager (CAM)
Research Team:	**York University:** Joydeep Mukherjee Marin Litoiu Cornel Barna

— *B2:* Modeling AUTOSAR Implementations in Simulink

Booth Number:	B2
Title:	Modeling AUTOSAR Implementations in Simulink

Give Feedback

Chair(s): Jian Chen
Format: Poster
Description: AUTOSAR defines the automotive three-layered software architecture. Application layer specifies functional behaviours and are realized as a set of tasks. Simulink is employed to address AUTOSAR's lack of support for modelling behaviours of tasks. Simulink simulations assume tasks are completed in zero execution time, while real executions require a finite execution time. We present a Simulink block that can schedule tasks with a non-zero simulation time to enable a more realistic simulation.
Industry Company: Queen's University
Product Name: Model Scheduler
Research Team: **Queen's University:** Jian Chen Thomas Dean
Ryerson University: Manar Alalfi
General Motors R&D: Ramesh S

— *B3:* Data Science 3 - Student Success Projects

Booth Number: B3
Title: Data Science 3 - Student Success Projects
Chair(s): Marcellus Mindel , Jennifer Collins , Dennis Buttera
Format: Poster
Description: In the third of three exhibits on Data Science Education, we will discuss how students at three different institutions are using Watson Studio to analyze institutional data to improve student retention and satisfaction. This approach is intended to create a pattern for supporting a wide variety of student capstone projects involving IBM tools and practices. Drop by to learn more about tools like Watson Studio and practices like Enterprise Design Thinking and Outcome Delivery.
Industry Company: IBM
Research Team: **IBM Advanced Studies:** Marcellus Mindel Jennifer Collins Dennis Buttera

— *B4:* Accessibility in the DevOps Era

Give Feedback

Booth Number:	B4
Title:	Accessibility in the DevOps Era
Chair(s):	Ali Asghar , Larry Lindsay , Thomas Brunet , Babz Kukoyi
Format:	Demo
Description:	As enterprises adopt continuous deployment practices, the DevOps testing practice must adapt accessibility testing tools and processes or user experience will suffer. These tools must work hand in hand with your application DevOps methodology - both required by & enabling those practices. In this hands-on demo, the audience will learn about IBM's cutting-edge accessibility DevOps tooling and see how simple it is to integrate the tools into an existing continuous integration pipeline (CI).
Web Site:	https://ibm.biz/a11yDashboard
Research Team:	**IBM:** Ali Jason Babz Kukoyi

— *B5:* Space Sentry Challenge: A game built with IoT devices and Microservices

Booth Number:	B5
Title:	Space Sentry Challenge: A game built with IoT devices and Microservices
Chair(s):	Frank Ji , Prashanth Gunapalasingam
Format:	Demo
Description:	Space Sentry Challenge is a fun arcade game that allows players to take control of a 3D printed spaceship to shoot down pop-up targets. This game is built for the purpose of showcasing the application of Microservices in an IoT context. It serves as a good sample project for anyone who wants to build IoT projects with Microservices. It uses popular IoT kits like Arduino, ESP wifi modules and Raspberry Pi with Microservices Java applications running on Open Liberty.
Business unit	IBM Cloud
Product Name:	Open Liberty
Web Site:	https://openliberty.io
Research Team:	**IBM:** Frank Ji Prashanth Gunapalasingam King Lam

Give Feedback

— *B6:* Comparative Analysis of Cloud Console Systems

Booth Number:	B6
Title:	Comparative Analysis of Cloud Console Systems
Chair(s):	Julia Rubin , Satish Sarraf , John Steinbacher , Anthony W Erwin
Format:	Poster
Description:	Cloud Console systems feature a customizable presentation layer that allows users to manage their cloud resources, such as VMs, storage, etc. Designing such systems is challenging due to the demand to (a) handle massive amounts of data, (b) integrate with different resource providers, and (c) deliver data from these providers in near real-time. In this work, we analyze several Cloud Consoles, including that of IBM Cloud, and discuss design decisions and trade-offs made to satisfy these demands.
Industry Company:	IBM
Project Number	1040
Product Name:	IBM Cloud
Research Team:	**University of British Columbia:** Julia Rubin Satish Sarraf Harshavardhan Kadiyala **Seneca College:** Mark Shtern **IBM Canada:** John Steinbacher **IBM USA:** Anthony W Erwin

— *B7:* REST APIs with LoopBack 4 and OpenAPI 3

Booth Number:	B7
Title:	REST APIs with LoopBack 4 and OpenAPI 3
Chair(s):	Taranveer Virk , Diana Lau , Biniam Admikew , Janny Hou
Format:	Poster
Description:	REST APIs power web applications, mobile applications, IoT devices & more. LoopBack 4 is an API creation framework that has been rewritten from the ground up and features a new programming interface, OpenAPI support, Dependency Injection, a new CLI and so much more. LoopBack makes it easy to create powerful OpenAPI 3.x (an industry standard) compliant REST APIs in minutes. It supports

Give Feedback

	connecting to multiple SQL and NoSQL databases. Poster will show how to get started in minutes with LoopBack 4.
Business unit	Hybrid Cloud
Product Name:	LoopBack
Web Site:	http://v4.loopback.io
Research Team:	**IBM:** Taranveer Virk Diana Lau Biniam Admikew Janny Hou

— *B8:* Interactive Visualization to Engage End-users in Prescriptive Analysis

Booth Number:	B8
Title:	Interactive Visualization to Engage End-users in Prescriptive Analysis
Chair(s):	Aindrila Basak
Format:	Poster
Description:	The research proposes a Visual Analytics System that can help end-users of Data Programming to visualize and understand patterns in input datasets and to write better labeling functions for generating labels. The proposed approach aims at enhancing the end-users' trust in predictive models and helping with the conversion of the predictions into actions. The design requirements of this work were identified by performing semi-structured interviews with end-users, data analysts, and domain experts.
Industry Company:	IBM
Research Team:	**University of Alberta:** Aindrila Basak Mona Nashaat James Miller **IBM Canada:** Shaikh Quader **IBM US:** Chad Marston

— *B9:* Finding Missing Labels for Large Industrial Datasets

Booth Number:	B9
Title:	Finding Missing Labels for Large Industrial Datasets
Chair(s):	Mona Nashaat
Format:	Poster

Give Feedback

Description: Obtaining labelled training datasets is a common challenge in the business domain. It costs a considerable amount of time and domain experience. We present a new approach to generate labelled datasets, in which we employ active learning in the data programming process. The proposed method aims at improving the accuracy of the training data. It optimizes the user engagement with the labelling process along with the annotation cost. It helps the users to trust the predictions of the final model.

Industry Company: IBM

Research Team: **University of Alberta:** Mona Nashaat Aindrila Basak James Miller
IBM Canada: Shaikh Quader
IBM US: Chad Marston

— *B10:* Synthesizing Realistic Data for Machine Learning

Booth Number: B10

Title: Synthesizing Realistic Data for Machine Learning

Chair(s): DAVID DCOSTA

Format: Poster

Description: This exhibit is to show several techniques to generate large volumes of synthetic data required by machine learning algorithms. The aspects explored will show the need and how to build proper relationships in attribute features using copulas to assist in testing machine learning algorithms. We will show ways to efficiently build and output large datasets in the order of 5 Million clients. In addition, a way to scale dataset creation in a parallel process across several virtual machines.

Product Name: Client Insights for Wealth Management

Web Site: https://www.ibm.com/us-en/marketplace/client-insight-for-wealth-management

Research Team: **IBM Canada Lab:** DAVID D'COSTA DIANE REYNOLDS DAVID XIE DONNA LEE PIA MEDINA CONNOR OBRIEN JOSH VALCHAR

— *B11:* Watch Phishing Evolution Over time

Booth Number: B11

Give Feedback

Title:	Watch Phishing Evolution Over time
Chair(s):	Qian Cui , Sophie Le Page , Emad Badawi , Ali Almokhtar
Format:	Poster
Description:	We proposed a model to monitor and analyze the phishing evolution. By using our model, we could observe the evolution pattern over different phishing groups
Industry Company:	IBM
Web Site:	https://esorics2018.upc.edu/program.do
Research Team:	**University of Ottawa:** Gregor v. Bochmann Guy-Vincent Jourdan Qian Cui Emad Badawi Ali Almokhtar Sophie Le Page **Polytechnique Montreal:** Ettore Merlo Stephane Heudron **IBM-CAS:** Iosif-Viorel (Vio) Onut

— *B12:* Easy Cloud-Native Development with Microclimate

Booth Number:	B12
Title:	Easy Cloud-Native Development with Microclimate
Chair(s):	Elson Yuen , John Collier , Sakib Hasan , Stephanie Cao
Format:	Poster
Description:	Microclimate is a brand-new development environment that offers a end-to-end development experience for Microservices from local or IBM Cloud private cluster. With Microclimate, you can create or import Java, Node.js, or Swift applications, and using any editor of your choosing to develop applications. We will demonstrate Microclimate's features on Iterative Development, DevOps pipeline and Intelligent Feedback, and will show how the tools will help developers to accelerate their development.
Business unit	Hybrid Cloud
Product Name:	Microclimate
Web Site:	https://microclimate-dev2ops.github.io/
Research Team:	**IBM:** Elson Yuen Eric Peters Erin Harris Jingfu Wang John Collier John Pitman Jonathan West Joseph Kim Laven Sathiyanathan Leonard Theivendra Keith Chong Maysun Faisal Rajiv Senthilnathan Sakib Haasan Stephanie Cao Steven Hung Tiaoyu Wang Tim Etchells

Give Feedback

— *B13:* A Machine Learning Approach to Test Case Prioritization

Booth Number:	B13
Title:	A Machine Learning Approach to Test Case Prioritization
Chair(s):	Francis Palma
Format:	Poster
Description:	We propose to use the defect-proneness value of source files as the essential measure for test case prioritization. Our proposed test case prioritization model comprises two stages: (1) Metric extraction and defect prediction and (2) Test case rescheduling.
Industry Company:	IBM
Project Number	Project SRG_2014_1
Product Name:	Test case prioritization
Research Team:	**Ryerson University:** Francis Palma and Ayse Bener

— *B14:* Optimization Driven Inlining

Booth Number:	B14
Title:	Optimization Driven Inlining
Chair(s):	Erick Ochoa
Format:	Poster
Description:	Inlining is a compiler optimization that should be selectively applied in JIT compilation. Inlining too much increases compilation time and inlining too little decreases throughput. Different measures (like invocation count and method size) have been used to decide which methods to inline. Our proposal is to use abstract interpretation to determine which optimizations are feasible and use this calculation to drive inlining.
Industry Company:	IBM
Project Number	1047
Product Name:	OpenJ9

Give Feedback

Research Team: **University of Alberta:** Erick Ochoa Karim Ali Jose Nelson Amaral
IBM Canada: Andrew Craik

— *B15:* Semantic-Aware Disambiguation for Entity Resolution

Booth Number: B15
Title: Semantic-Aware Disambiguation for Entity Resolution
Chair(s): Yu Huang
Format: Poster
Description: As data is integrated, duplicates inevitably occur in the integrated instance. Existing similarity functions often rely on string similarity to identify duplicates, which is insufficient in capturing the semantic meaning of terms. We propose a semantic-aware deduplication framework which can identify the duplicates by extracting the semantic information to measure 'closeness' of entities.
Industry Company: IBM
Research Team: **McMaster University:** Yu Huang Fei Chiang
IBM Germany: Albert Maier Yannick Saillet
IBM Chicago: Damir Spisic
IBM Toronto Lab: Calisto Zuzarte

— *B16:* Create data-driven visual reports in Reactive Reports, a framework built for big data

Booth Number: B16
Title: Create data-driven visual reports in Reactive Reports, a framework built for big data
Chair(s): Tereza Nedelescu , Mihai Nicolae , Samar Sajnani
Format: Demo
Description: This expo will showcase Reactive Reports, an open source framework for generating data-driven reports. The framework can use Scala and Java code to take input from mixed sources, such as database records, objects, or streams. The generated output is one or more PDF reports that can be used to summarize or

Give Feedback

visualize data. Reactive Reports scales well for large amounts of input data due to its memory efficient design, the use of Akka streaming, and its internal data storage architecture.

Industry Company: N/A

Product Name: Reactive Reports

Web Site: https://github.com/sysalto/ReactiveReports

Research Team: **IBM Canada:** Tereza Nedelescu Mihai Nicolae Samar Sajnani

— *B17:* OptImatch: Semantic-Web System for Automated Query Tuning

Booth Number: B17

Title: OptImatch: Semantic-Web System for Automated Query Tuning

Chair(s): Guilherme Damasio , alexandar mihaylov

Format: Poster

Description: Optimatch is a system which discovers problem patterns automatically by matching them to templates of plans from a previously populated knowledge base. The system extends upon our previous work, but now alleviating all need for manual expert input by introducing a customizable knowledge base capable of learning problem patterns from previously ran execution plans. The system output is a query with the embedded optimization guideline which enforces the optimal execution plan to be applied.

Industry Company: IBM

Project Number 1049

Product Name: Optimatch

Research Team: **University of Ontario Institute of Tech.:** Guilherme Fetter Damasio Alexandar Mihaylov Jaroslaw Szlichta

York University: Parke Godfrey

IBM: Calisto Zuzarte Piotr Mierzejewski Vincent Corvinelli

— *B18:* Third Party Container Tests

Booth Number: B18

Give Feedback

Title:	Third Party Container Tests
Chair(s):	Mesbah Alam
Format:	Poster
Description:	This poster session depicts the '3rd party container tests' that we curate and run at the AdoptOpenJDK project. The IBM Runtime technologies team that produces the 'Eclipse OpenJ9' product uses many test materials to verify the SDK binaries. In case of 3rd party application tests, we run a variety of Java applications inside Docker containers. This poster session will describe the test automation techniques we use to execute, triage and debug application tests inside Docker.
Business unit	IBM Runtime Technologies
Product Name:	OpenJDK with Eclipse OpenJ9
Research Team:	**IBM:** Mesbah

— *B19:* Runtime Automatic Parallelization of JVM Applications

Booth Number:	B19
Title:	Runtime Automatic Parallelization of JVM Applications
Chair(s):	Akihiro Hayashi , Gita Koblents
Format:	Poster
Description:	In this IBM-Rice University collaboration, we push the bounds on integrating JVM applications with GPU accelerators. This work uses advanced code generation and performance prediction techniques to effectively and automatically offload JVM applications including Apache Spark to hardware accelerators.
Industry Company:	IBM
Project Number	900
Product Name:	OpenJ9, Spark-DK
Research Team:	**Rice University:** Akihiro Hayashi **IBM Canada:** Gita Koblents Jimmy Kwa **Georgia Institute of Technology:** Vivek Sarkar **IBM Research Tokyo:** Kazuaki Ishizaki

Give Feedback

— *B20:* Variability Aware analysis using Clang

Booth Number:	B20
Title:	Variability Aware analysis using Clang
Chair(s):	Jacob Reckhard
Format:	Poster
Description:	Tool-based analysis is important to catch errors before the codebase becomes too dependent on the erroneous parts. One well known project for compilation and analysis of code is Clang. Here, we describe the changes required to coerce Clang into conducting variability-aware analysis of C++ code. Varibility-aware analysis is a technique in which the code base is analysed for all possible configurations of the code.
Industry Company:	IBM
Research Team:	**University of Alberta:** Jacob Reckhard Dr. Sarah Nadi

— *B21:* Approximation Algorithms for Instruction Scheduling and Rapid Prototyping in Coconut

Booth Number:	B21
Title:	Approximation Algorithms for Instruction Scheduling and Rapid Prototyping in Coconut
Chair(s):	Curtis D'Alves , Bill O'Farrell
Format:	Poster
Description:	Instruction scheduling is an NP-complete problem which allows for altering the execution order of instructions in a function without altering the function’s semantics. We present a continuous optimization based model for solving near-optimal schedules and a prototype implemented for the IBM z13 and z14 architectures and used to generate fast schedules for common cryptographic functions used in Blockchain operations implemented in GOLang assembly.
Industry Company:	IBM
Project Number	1006
Product Name:	Blockchain Cryptographic Functions in GO Assembly for Z
Research Team:	**McMaster University:** Curtis D'Alves Bill O'Farrell Christopher Anand Wolfram Kahl James You

Give Feedback

— *B22:* Static Dataflow Analysis for OpenMP Code Sanitization

Booth Number:	B22
Title:	Static Dataflow Analysis for OpenMP Code Sanitization
Chair(s):	Prithayan Barua , Jun Shirako , Whitney Tsang , Vivek Sarkar
Format:	Poster
Description:	OpenMP provides high-level abstractions and rich functionality for GPU programming. However, understanding the complex semantics of various combinations of pragmas is not trivial and leads to their wrong usage. We present a static analysis technique to detect incorrect usage of OpenMP target data mapping. Assuming the correctness of underlying sequential programs, our approach detects the OpenMP errors by comparing the data flow information of the original OpenMP version with the sequential.
Industry Company:	IBM
Project Number	CAS-896
Product Name:	XL Compiler
Research Team:	**Georgia Institute of Technology:** Prithayan Barua, Jun Shirako, Vivek Sarkar **IBM Toronto:** Whitney T Tsang, Jeeva Paudel, Wang Chen

— *B23:* Plasma Focused Ion Beam Curtaining Artefact Correction by Fourier-Based Linear Model

Booth Number:	B23
Title:	Plasma Focused Ion Beam Curtaining Artefact Correction by Fourier-Based Linear Model
Chair(s):	Christopher Schankula , Christopher Anand
Format:	Poster
Description:	This exhibit describes one example of a multidisciplinary partnership between the McMaster Computing & Software and Materials Science & Engineering departments to better solve modern Materials Science problems computationally. Specifically, we describe a software package for reducing artefacts in focused ion beam scanning electron microscope (FIB-SEM) images. We will describe and provide a live demonstration of the linear optimization model used to reduce the curtaining artefacts.

Give Feedback

Industry Company:	IBM
Research Team:	**McMaster University:** Christopher W. Schankula Christopher K. Anand Nabil D. Bassim

— *B24:* FLASH: Multithreaded Fault Localization

Booth Number:	B24
Title:	FLASH: Multithreaded Fault Localization
Chair(s):	Luisa Rojas Garcia , Jeremy Bradbury
Format:	Poster
Description:	Multithreaded programs can have considerable performance benefits over sequential programs. However, this often comes at a cost with respect to program understandability as well as testing and debugging. To address these challenges we have developed an automatic fault localization tool for multithreaded Java programs (FLASH). Our tool utilizes a combination of code injection and a heuristic search algorithm to identify high probability code blocks.
Industry Company:	IBM
Research Team:	**University of Ontario IT:** Luisa Rojas Garcia Jeremy Bradbury

Give Feedback

People

Think: Cognitive Computing, Big Data, Cloud, Security and Privacy

October 29 - 31

CASCON News

Download the CASCON 2018 Program Guide

See the CASCON 2018 Proceeding in ACM

Get certain Workshops materials

Call for Papers for CASCON 2019

Give Feedback

Organizing Committee

Name	Position	Institution	Social Profile

Name	Position	Institution	Social Profile
Vio Onut	Conference Chair	IBM Canada Ltd	
Dorina Petriu	Conference Program Co-Chair	Carleton University	
Wang Chen	Conference Program Co-Chair	IBM Canada Ltd	
Hausi Müller	Workshop Co-Chair	University of Victoria	
Tinny Ng	Workshop Co-Chair	IBM Canada Ltd	
Kenneth Kent	Expo Co-Chair	University of New Brunswick	
Robert Enenkel	Expo Co-Chair	IBM Canada Ltd	
Marcellus Mindel	Finance and Registration Chair	IBM Canada Ltd	
Harnoor Manan	Website Chair	IBM Canada Ltd	
Marin Litoiu	Local Arrangements Co-Chair	York University	
Vio Onut	Local Arrangements Co-Chair	IBM Canada Ltd	

Give Feedback

Name	Position	Institution	Social Profile
Marin Litoiu	Volunteer Chair	York University	
Guy-Vincent Jourdan	Publication Chair	University of Ottawa	
Andrew Jaramillo	Conference Proceedings Editor	IBM Canada Ltd	
Rodney D'Silva	Network Chair	IBM Canada Ltd	
Alan Heighway	Network Chair	IBM Canada Ltd	

Steering Committee

Name	Position	Institution	Social Profile
Marin Litoiu	Steering Committee	York University	
Kelly Lyons	Steering Committee	University of Toronto	
Marcellus Mindel	Steering Committee	IBM Canada Ltd	
Hausi Müller	Steering Committee	University of Victoria	

Give Feedback

Name	Position	Institution	Social Profile
Tinny Ng	Steering Committee	IBM Canada Ltd	
Vio Onut	Steering Committee	IBM Canada Ltd	
Joe Wigglesworth	Steering Committee	IBM Canada Ltd	
Ken Wong	Steering Committee	University of Alberta	
Ying Zou	Steering Committee	Queen's University	

Program Committee

Name	Position	Institution	Social Profile
Bram Adams	Program Committee	École Polytechnique de Montréal	
Jose Nelson Amaral	Program Committee	University of Alberta	
Daniel Amyot	Program Committee	University of Ottawa	
Christopher Anand	Program Committee	McMaster University	

Give Feedback

Name	Position	Institution	Social Profile
Ebrahim Bagheri	Program Committee	Ryerson University	
Ayse Bener	Program Committee	Ryerson University	
Jeremy Bradbury	Program Committee	University of Ontario Institute of Technology	
David Bremner	Program Committee	University of New Brunswick	
Tse-Hsun Peter Chen	Program Committee	Concordia University	
Mark Chignell	Program Committee	University of Toronto	
Andrew Craik	Program Committee	IBM Canada Ltd	
Frank Dehne	Program Committee	Carleton University	
Chen Ding	Program Committee	University of Rochester	
Juergen Dingel	Program Committee	Queen's University	
Herhard Dueck	Program Committee	University of New Brunswick	

Give Feedback

Name	Position	Institution	Social Profile
Ghizlane El Boussaidi	Program Committee	École de Technologie Supérieure	
Michalis Famelis	Program Committee	University of Montreal	
Marios Fokaefs	Program Committee	École Polytechnique de Montréal	
Sudhakar Ganti	Program Committee	University of Victoria	
James Green	Program Committee	Carleton University	
Arie Gurfinkel	Program Committee	University of Waterloo	
Biruk Habtemariam	Program Committee	IBM Canada Ltd	
Ahmed Hassan	Program Committee	Queen's University	
Akihiro Hayashi	Program Committee	Rice University	
Hadi Hemmati	Program Committee	University of Calgary	
Hans-Arno Jacobsen	Program Committee	University of Toronto	

Give Feedback

Name	Position	Institution	Social Profile
Guy-Vincent Jourdan	Program Committee	University of Ottawa	
Wolfram Kahl	Program Committee	McMaster University	
Kenneth Kent	Program Committee	University of New Brunswick	
Foutse Khomh	Program Committee	École Polytechnique de Montréal	
Kostas Kontogiannis	Program Committee	University of Western Ontario	
Diwakar Krishnamurthy	Program Committee	University of Calgary	
Alexei Lapouchnian	Program Committee	University of Toronto	
Grace Lewis	Program Committee	Carnegie Mellon Software Engineering Institute	
Jin Li	Program Committee	PointClickCare	
Marin Litoiu	Program Committee	York University	
Hanan Lutfiyya	Program Committee	University of Western Ontario	

Give Feedback

Name	Position	Institution	Social Profile
Nazim Madhavji	Program Committee	University of Western Ontario	
Daryl Maier	Program Committee	IBM Canada Ltd	
Pat Martin	Program Committee	Queen's University	
Stan Matwin	Program Committee	Dalhousie University	
Ettore Merlo	Program Committee	École Polytechnique de Montréal	
Andriy Miranskyy	Program Committee	Ryerson University	
Marc Moreno Maza	Program Committee	University of Western Ontario	
Hausi Müller	Program Committee	University of Victoria	
Maleknaz Nayebi	Program Committee	University of Toronto	
Tinny Ng	Program Committee	IBM Canada Ltd	
Bill O'Farrell	Program Committee	IBM Canada Ltd	

Give Feedback

Name	Position	Institution	Social Profile
Ipek Ozkaya	Program Committee	Carnegie Mellon Software Engineering Institute	
Gennady Pekhimenko	Program Committee	Carnegie Mellon University	
Fred Popowich	Program Committee	Simon Fraser University	
Suprio Ray	Program Committee	University of New Brunswick	
Juergen Rilling	Program Committee	Concordia University	
Walid Rjaibi	Program Committee	IBM Canada Ltd	
Chanchal K. Roy	Program Committee	University of Saskatchewan	
Weiyi Shang	Program Committee	Concordia University	
Xipeng Shen	Program Committee	North Carolina State University	
Jun Shirak	Program Committee	Georgia Institute of Technology	
Gabriel Silberman	Program Committee	Barcelona Institute of Science and Technology	

Give Feedback

Name	Position	Institution	Social Profile
Michael Smit	Program Committee	Dalhousie University	
Ulrike Stege	Program Committee	University of Victoria	
Mark Stoodley	Program Committee	IBM Canada Ltd	
Jaroslaw Szlichta	Program Committee	University of Ontario Institute of Technology	
Ladan Tahvildari	Program Committee	University of Waterloo	
Norha M. Villega	Program Committee	Universidad Icesi, Cali, Colombia	
Paul Ward	Program Committee	University of Waterloo	
Joe Wigglesworth	Program Committee	IBM Canada Ltd	
Eric Yu	Program Committee	University of Toronto	
Ying Zou	Program Committee	Queen's University	
Farhana Zulkernine	Program Committee	Queen's University	

Give Feedback

Name	Position	Institution	Social Profile
Calisto Zuzarte	Program Committee	IBM Canada Ltd	

Best Paper Awards Selection Committee

Name	Position	Institution	Social Profile
Dorina Petriu	Best Paper Awards Selection Committee Co-Chair	Carleton University	
Wang Chen	Best Paper Awards Selection Committee Co-Chair	IBM Canada Ltd	
Ebrahim Bagheri	Best Paper Awards Selection Committee	Ryerson University	
Andrew Craik	Best Paper Awards Selection Committee	IBM Canada Ltd	
Juergen Dingel	Best Paper Awards Selection Committee	Queen's University	

Most Influential Paper Award Committee

Name	Position	Institution	Social Profile

Give Feedback

Name	Position	Institution	Social Profile
Dorina Petriu	Most Influential Paper Award Committee Co-Chair	Carleton University	
Wang Chen	Most Influential Paper Award Committee Co-Chair	IBM Canada Ltd	
Marsha Chechik	Most Influential Paper Award Committee	University of Toronto	
Ken Wong	Most Influential Paper Award Committee	University of Alberta	

Workshop Selection Committee

Name	Position	Institution	Social Profile
Hausi Müller	Workshop Co-Chair	University of Victoria	
Tinny Ng	Workshop Co-Chair	IBM Canada Ltd	
Frank Li	Workshop Selection Committee	IBM Canada Ltd	
Marin Litoiu	Workshop Selection Committee	York University	
Kishor Patil	Workshop Selection Committee	IBM Canada Ltd	

Give Feedback

Name	Position	Institution	Social Profile
Maria Patrou	Workshop Selection Committee	University of New Brunswick	
Ladan Tahvildari	Workshop Selection Committee	University of Waterloo	
Derek Truong	Workshop Selection Committee	IBM Canada Ltd	
Ying Zou	Workshop Selection Committee	Queen's University	
Calisto Zuzarte	Workshop Selection Committee	IBM Canada Ltd	

Best Expo Award Selection Committee

Name	Position	Institution	Social Profile
Robert Enenkel	Expo Co-Chair	IBM Canada Ltd	
Kenneth Kent	Expo Co-Chair	University of New Brunswick	
Allan Kielstra	Best Expo Award Selection Committee	IBM Canada Ltd	
Allen Chan	Best Expo Award Selection Committee	IBM Canada Ltd	

Give Feedback

Name	Position	Institution	Social Profile
Danny Gupta	Best Expo Award Selection Committee	IBM Canada Ltd	
David Kalmuk	Best Expo Award Selection Committee	IBM Canada Ltd	
Joe Wigglesworth	Best Expo Award Selection Committee	IBM Canada Ltd	
John Steinbacher	Best Expo Award Selection Committee	IBM Canada Ltd	
Mark Stoodley	Best Expo Award Selection Committee	IBM Canada Ltd	
Piotr Mierzejewsk	Best Expo Award Selection Committee	IBM Canada Ltd	
Roch Archambault	Best Expo Award Selection Committee	IBM Canada Ltd	
Roland Koo	Best Expo Award Selection Committee	IBM Canada Ltd	
Vio Onut	Best Expo Award Selection Committee	IBM Canada Ltd	

Give Feedback

Venue

Think: Cognitive Computing, Big Data, Cloud, Security and Privacy

October 29 - 31

CASCON News

Download the CASCON 2018 Program Guide

See the CASCON 2018 Proceeding in ACM

Get certain Workshops materials

Call for Papers for CASCON 2019

Give Feedback

2018 CASCON Venue

October 29 - 31

Hilton Toronto / Markham Suites Conference Center and Spa

8500 Warden Avenue, Markham, Canada L6G 1A5

905-470-8500

Group Code: 1BMC

Note: The group code above can only be used to book a room from October 27th to November 1st.

The procedure to book a hotel room is as follows:

Step:1 Visit the **Hilton** website to start the booking process.

Step:2 Click the *"more options"* button shown below.

Step:3 Click the *"add special rates"* button and then add the group code *"1BMC"*.

NOTE: The room will not show if the date entered is not in the range of Oct 27 - Nov 1. Make sure you enter the code in "group code" field as the IBM CASCON 2018 room will not show if you enter it in the promo code field.

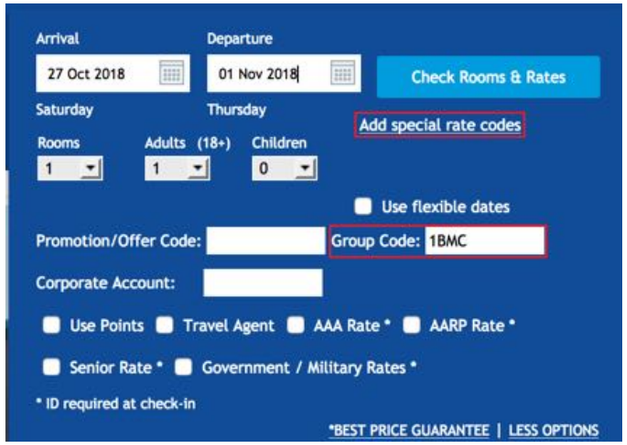

Give Feedback

Step:4 Click *"Check Rooms & Rates"* button. After redirect, in the new page you should see room details *IBM CASCON 2018*. There are more price options when you expand the section clicking the *"more prices"* button.

Step:5 Choose any other price options if needed and then click *"select"* you will be redirected to a new page where you fill out the form there are more options such as add another guest, special requests and free wifi which you can choose from. Click *Continue*, A new page opens where you can check the room details then you fill the payment information and click on *Book Reservation*.

Give Feedback

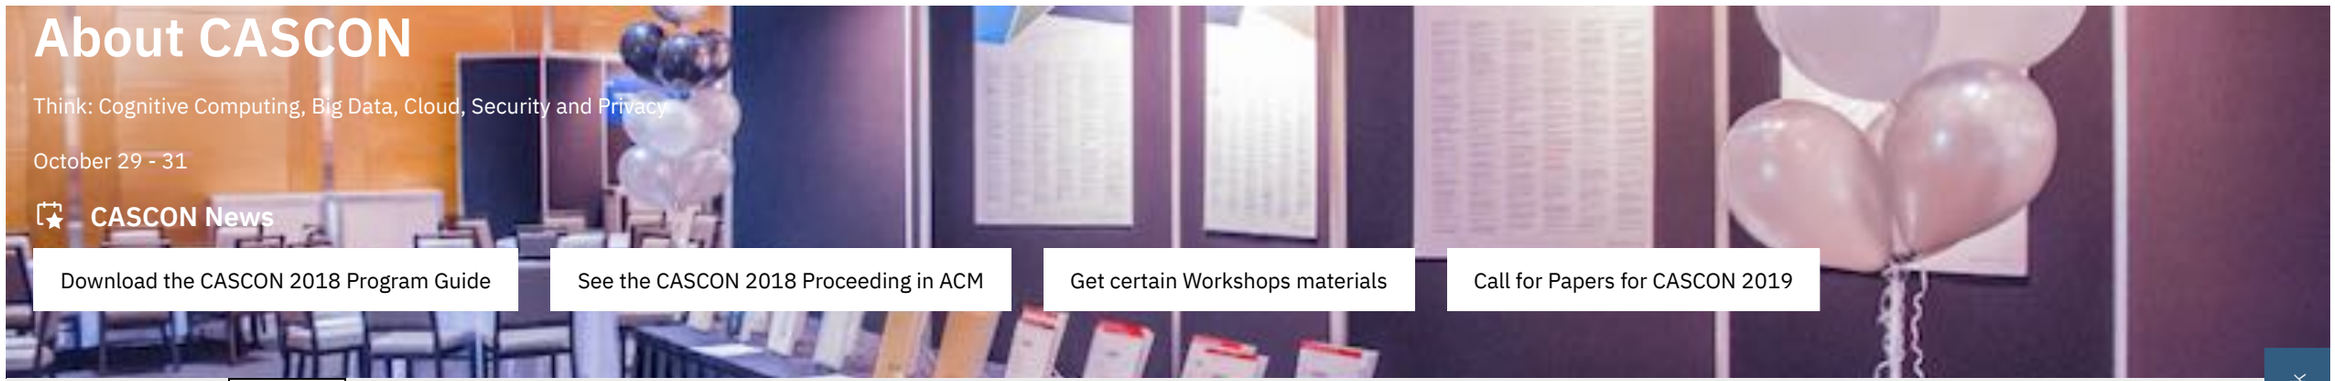

About CASCON

Think: Cognitive Computing, Big Data, Cloud, Security and Privacy

October 29 - 31

CASCON News

Download the CASCON 2018 Program Guide

See the CASCON 2018 Proceeding in ACM

Get certain Workshops materials

Call for Papers for CASCON 2019

CASCON is a premier industrial and academic conference in computer science and software engineering. Hosted by IBM Advanced Studies and IBM Canada Laboratory, CASCON attracts developers, researchers, innovators, technologists and decision makers from a variety of sectors including academia, industry, and government. We all come together to discover and discuss technology trends, share knowledge, form collaborations, write proposals, present papers, exhibit prototypes and showcase products.

CASCON 2018 Statistics

28
YEARS OF IMPACT

1036
ATTENDEES

4
KEYNOTES

33
PAPERS

81
EXPO

29
WORKSHOPS

Steering Committee Members

Name	Position	Institution	Social Profile
Marin Litoiu	Steering Committee	York University	
Kelly Lyons	Steering Committee	University of Toronto	
Marcellus Mindel	Steering Committee	IBM Canada Ltd	
Hausi Müller	Steering Committee	University of Victoria	
Tinny Ng	Steering Committee	IBM Canada Ltd	
Vio Onut	Steering Committee	IBM Canada Ltd	
Joe Wigglesworth	Steering Committee	IBM Canada Ltd	
Ken Wong	Steering Committee	University of Alberta	
Ying Zou	Steering Committee	Queen's University	

Give Feedback

Sponsors

Give Feedback

Give Feedback

IBM Advanced Studies(CAS)

https://www.ibm.com/ibm/cas/canada/

IBM Advanced Studies (CAS) is a system of engagement for academic collaboration with IBM. We network IBMers with a community of students, researchers, educators, change makers and entrepreneurs to attract talent, partners and clients to IBM. Together, we explore technical and social challenges around emerging architectures that will define business and computing for the next quarter century.

Give Feedback

Awards

Think: Cognitive Computing, Big Data, Cloud, Security and Privacy

October 29 - 31

CASCON News

Download the CASCON 2018 Program Guide

See the CASCON 2018 Proceeding in ACM

Get certain Workshops materials

Call for Papers for CASCON 2019

Give Feedback

Share the page:

Project of the Year

Steven Astorino, **Tanya Bouman, Lucas Dutton, Dr. Robert F. Enenkel, Professor Christopher Anand, Adele Olejarz, Yumna Irfan, Wang Chen**, Marcellus Mindel
(**Other winners**: Curtis D'Alves, Yusra Irfan, Professor Wolfram Kahl, Jinsong Ji, Marcel Mitran)

→ Click to read more and share

Faculty Fellow of the Year

Guy-Vincent Jourdan, Steven Astorino, Marcellus Mindel

→ Click to read more and share

Student of the Year

Steven Astorino, **Artem Chikin**, Marcellus Mindel

→ Click to read more and share

Collaborator of the Year

Best Paper

Best Student Paper

Give Feedback

Steven Astorino, **Andrew Craik**, Marcellus Mindel

→ Click to read more and share

Wang Chen, **David Wille, Kenny Wehling**, Dorina Petriu
(**Other winners**: Christoph Seidl and Ina Schaefer)

→ Click to read more and share

Dorina Petriu, Ettore Merlo, **Amine Barrak , Le An**, Foutse Khomh, Wang Chen
(**Other winner**: Marc-André Laverdière)

→ Click to read more and share

Most Influential Paper

Wang Chen, **Giuliano Antoniol, Foutse Khomh, Yann-Gaël Guéhéneuc**, Dorina Petriu

Best Expo Exhibit

Akihiro Hayashi, Kenneth Kent, Robert Enenkel
(**Other winner**: Gita Koblents)

Expo People's Choice Exhibit

Kenneth Kent, **Diana Lau, Janny Hou, Biniam Admikew**, Robert

Give Feedback

(**Other winners**: Kamel Ayari, Massimiliano Di Penta)

→ Click to read more and share

→ Click to read more and share

Enenkel
(**Other winner**: Taranveer Virk)

→ Click to read more and share

Give Feedback

Awards

Think: Cognitive Computing, Big Data, Cloud, Security and Privacy

October 29 - 31

CASCON News

Download the CASCON 2018 Program Guide

See the CASCON 2018 Proceeding in ACM

Get certain Workshops materials

Call for Papers for CASCON 2019

Give Feedback

Project of the Year

Faculty Fellow of the Year

Student of the Year

Tanya Woods, **Professor Guy-Vincent Jourdan, Sophie Le Page, Qian Cui, Prof. Gregor von Bochmann, Dr. Iosif Viorel Onut**, Marcellus Mindel

→ Click to read more and share

Marcellus Mindel, **Dr. Christopher Anand**, Nora Young

→ Click to read more and share

Marcellus Mindel, **Yu Zhao**, Nora Young

→ Click to read more and share

Collaborator of the Year

Marcellus Mindel, Kishor Patil(on behalf of **Michael Dawson**), Nora Young;

Best Paper

Joe Wigglesworth, Victor R Basili, **Xiang Jiang, Stan Matwin**, Kelly Lyons

Best Student Paper

Joe Wigglesworth, Victor R Basili, **Krystalenia Tatsi**, Kostas Kontogiannis, Kelly Lyons

Give Feedback

→ Click to read more and share

→ Click to read more and share

→ Click to read more and share

Most Influential Paper

Donna Dillenberger, **Bill O'Farrell, Jon Bennett, Yuan Guan**, Marcellus Mindel

→ Click to read more and share

Best Expo Exhibit

Amente Bekele, Dr.James R.Green, Marcellus Mindel, **Dr.Shermeen Nizami, Yasmina Souley Dosso; Other winner(s):** Akihiro Hayashi

→ Click to read more and share

Expo People's Choice Exhibit

Maria Patrou, Susan Landau, Marcellus Mindel, Kenneth B. Kent, Gerhard W. Dueck

→ Click to read more and share

Give Feedback

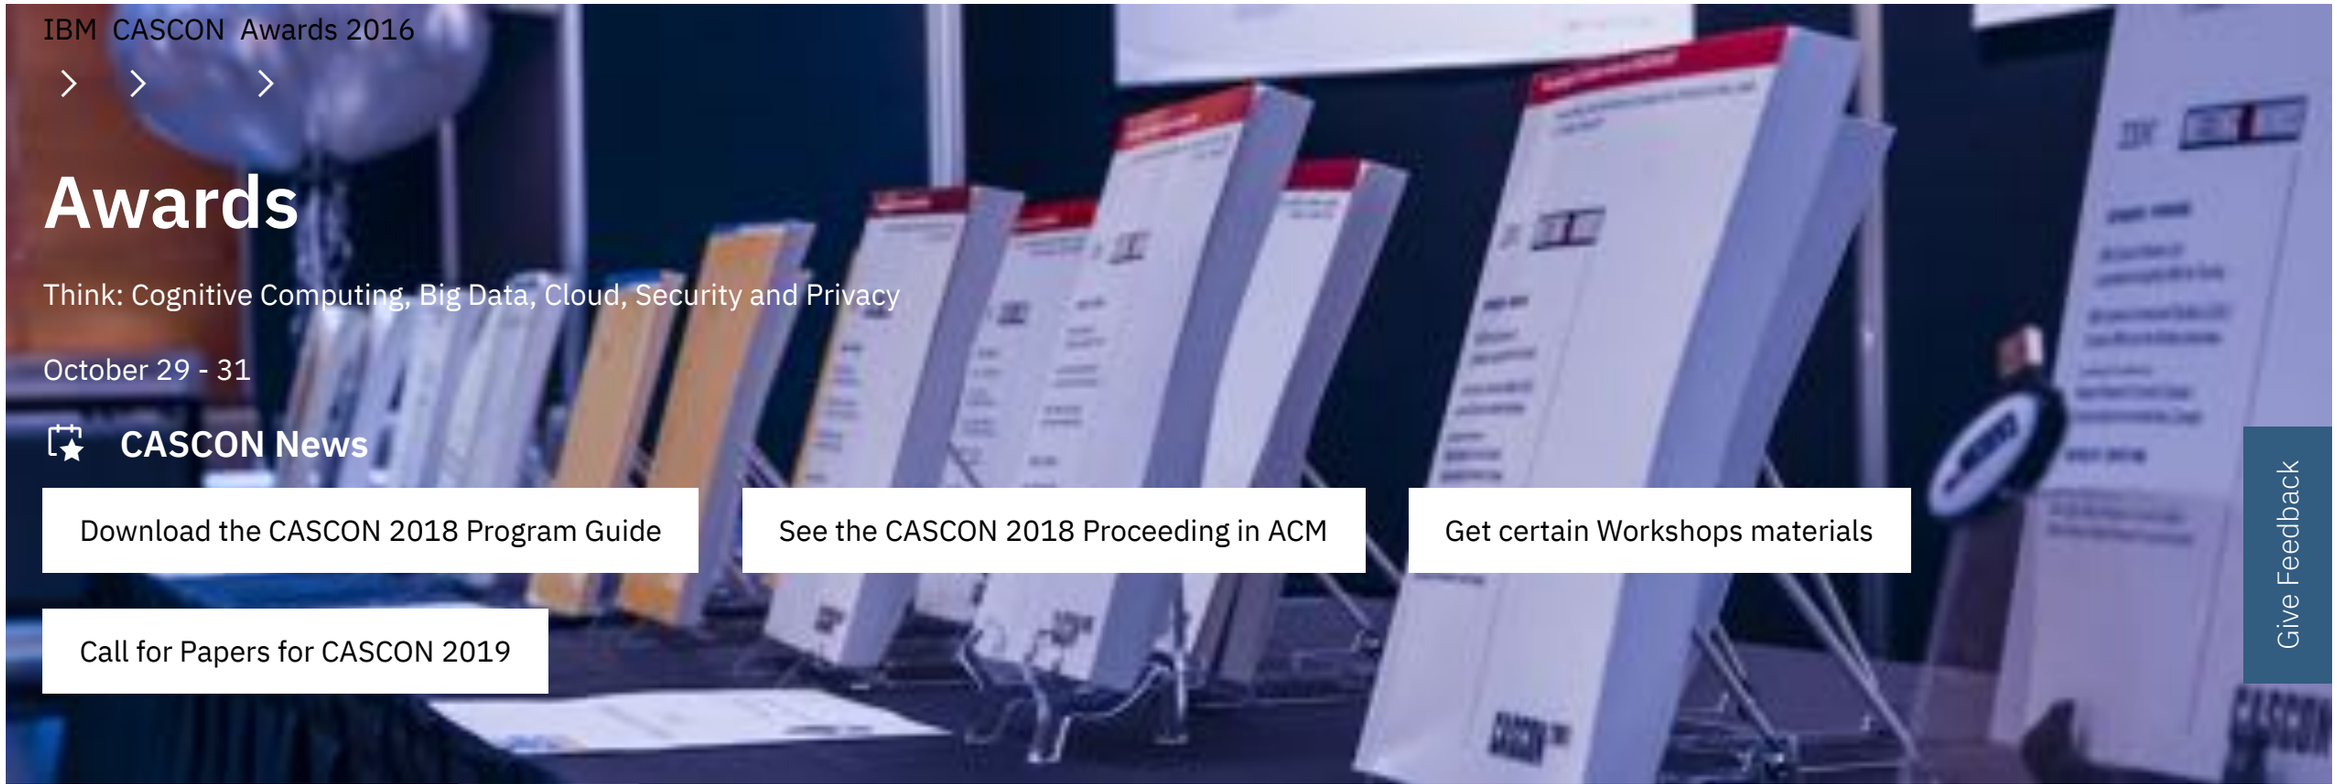

Awards

Think: Cognitive Computing, Big Data, Cloud, Security and Privacy

October 29 - 31

CASCON News

Download the CASCON 2018 Program Guide

See the CASCON 2018 Proceeding in ACM

Get certain Workshops materials

Call for Papers for CASCON 2019

Give Feedback

Project of the Year

Faculty Fellow of the Year

Student of the Year

→ Click to read more and share

→ Click to read more and share

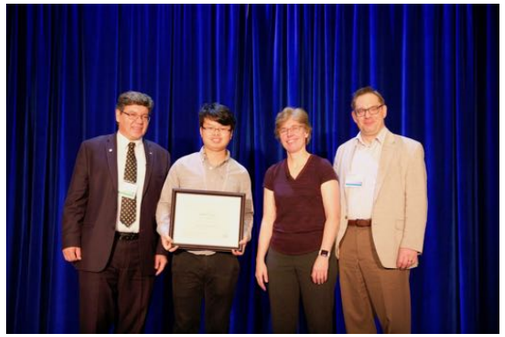

→ Click to read more and share

Collaborator of the Year

→ Click to read more and share

Innovator of the Year

→ Click to read more and share

Best Paper

→ Click to read more and share

Give Feedback

Best Student Paper

→ Click to read more and share

Most Influential Paper

→ Click to read more and share

Best Expo Exhibit

→ Click to read more and share

Expo People's Choice Exhibit of the Year

Give Feedback

Click to read more and share

Give Feedback

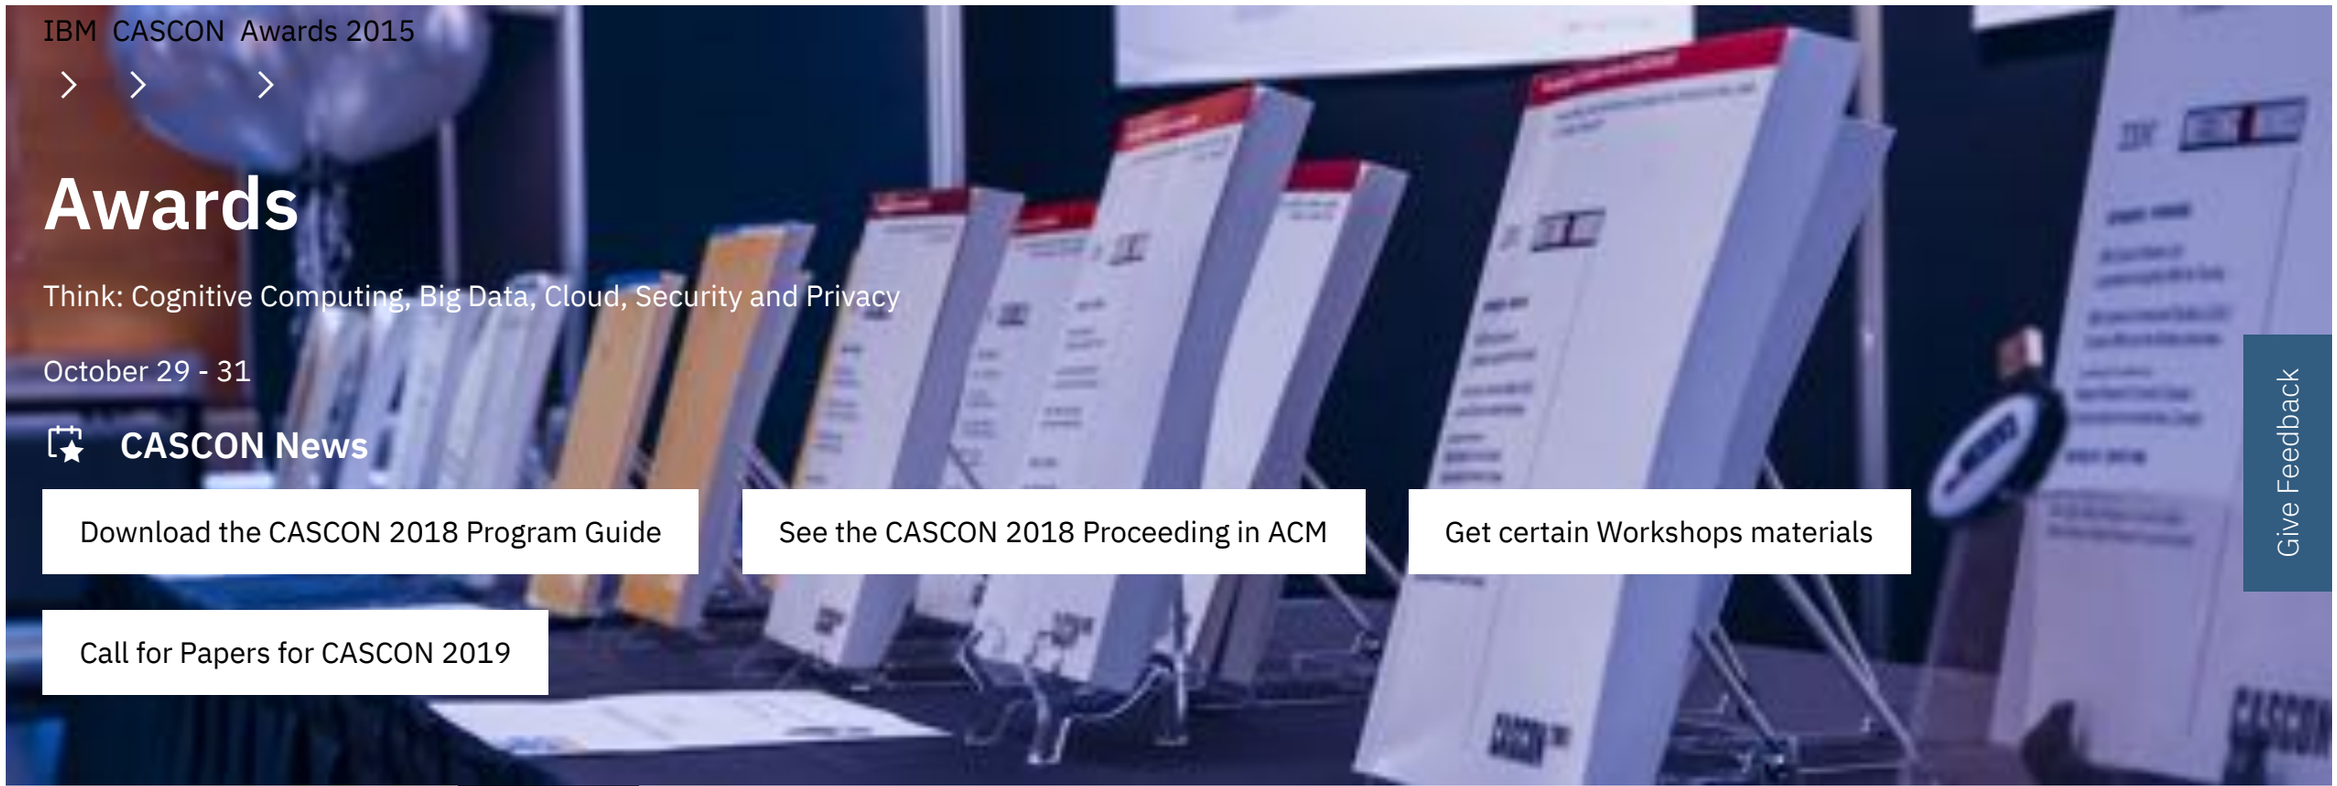

IBM CASCON Awards 2015

Awards

Think: Cognitive Computing, Big Data, Cloud, Security and Privacy

October 29 - 31

CASCON News

Download the CASCON 2018 Program Guide

See the CASCON 2018 Proceeding in ACM

Get certain Workshops materials

Call for Papers for CASCON 2019

Give Feedback

Collaborator of the Year

Faculty Fellow of the Year

Innovator of the Year

→ Click to read more and share

→ Click to read more and share

→ Click to read more and share

Student of the Year

→ Click to read more and share

Project of the Year

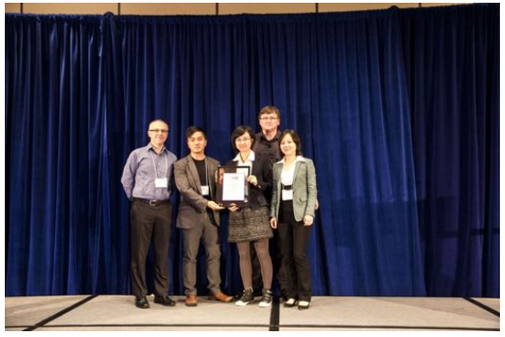

→ Click to read more and share

Best Paper

→ Click to read more and share

Give Feedback

Best Student Paper

→ Click to read more and share

Best Technology Showcase

→ Click to read more and share

Most Influential Paper

→ Click to read more and share

People's Choice Technology Showcase

Give Feedback

→ Click to read more and share

Give Feedback

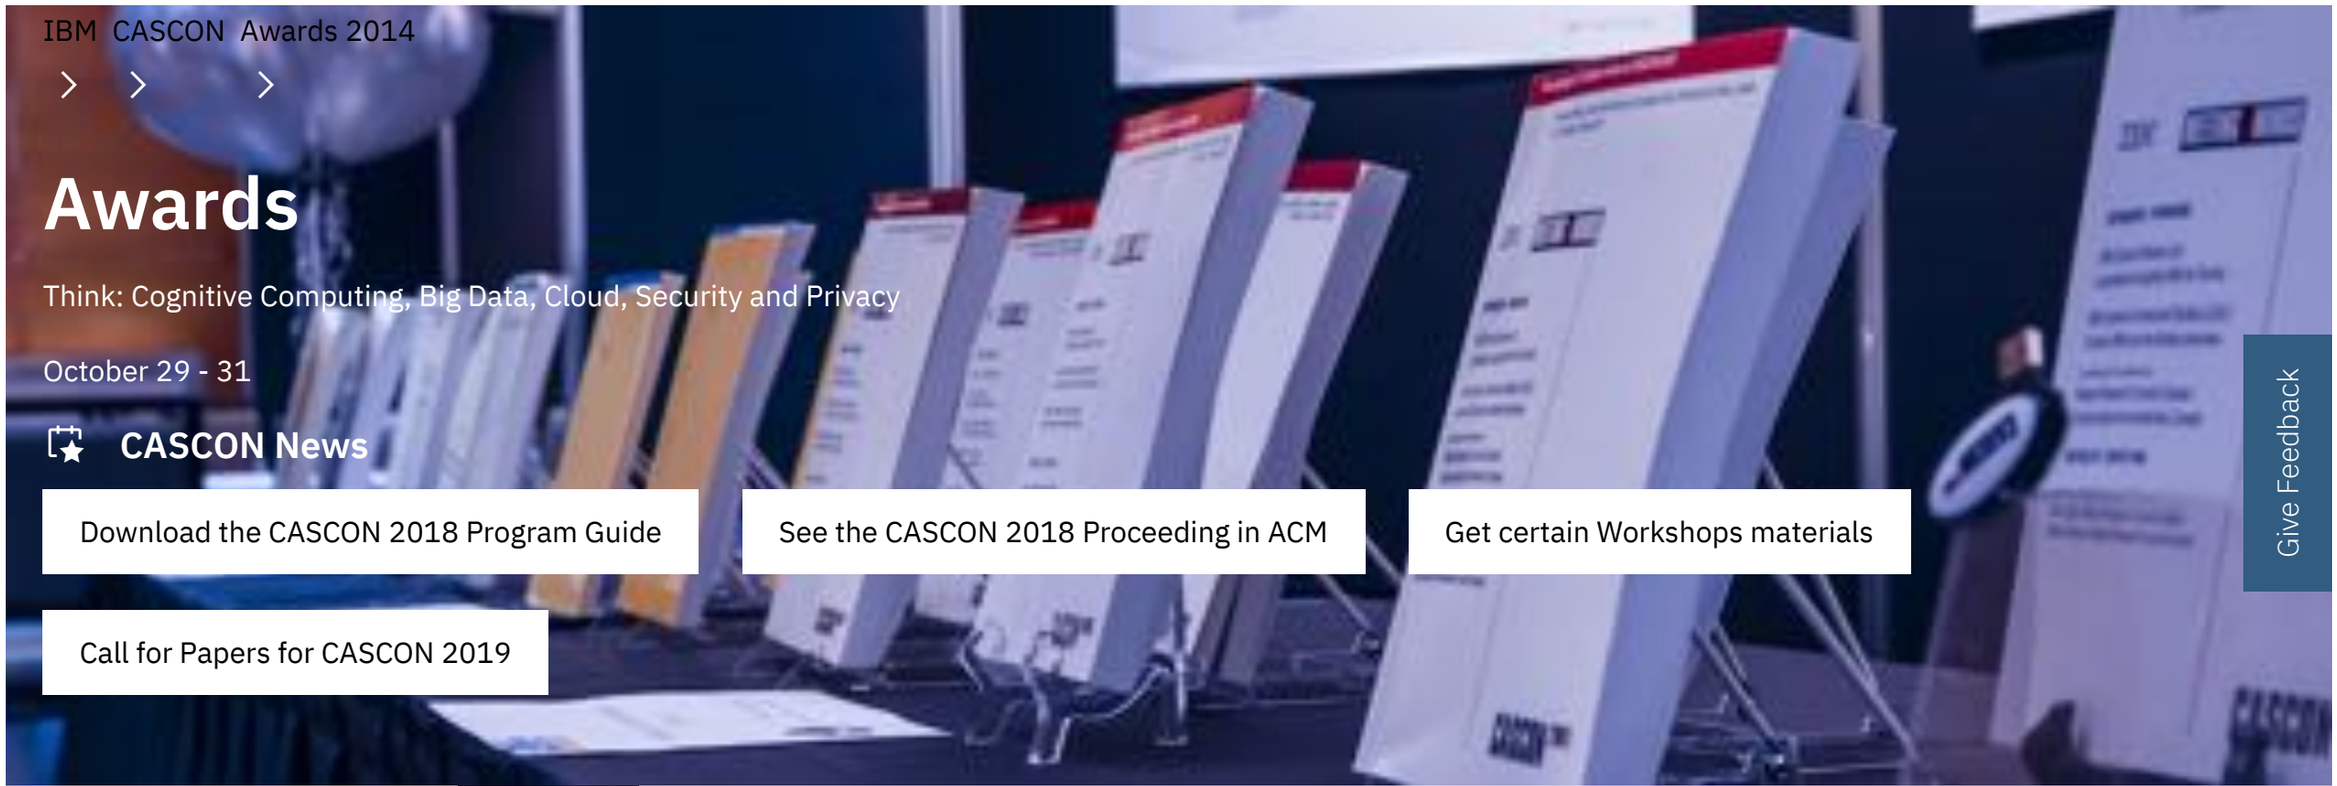

IBM CASCON Awards 2014

Awards

Think: Cognitive Computing, Big Data, Cloud, Security and Privacy

October 29 - 31

CASCON News

Download the CASCON 2018 Program Guide

See the CASCON 2018 Proceeding in ACM

Get certain Workshops materials

Call for Papers for CASCON 2019

Give Feedback

Collaborator of the Year

Faculty Fellow of the Year

Innovator of the Year

→ Click to read more and share

→ Click to read more and share

→ Click to read more and share

Project of the Year

→ Click to read more and share

Best Paper

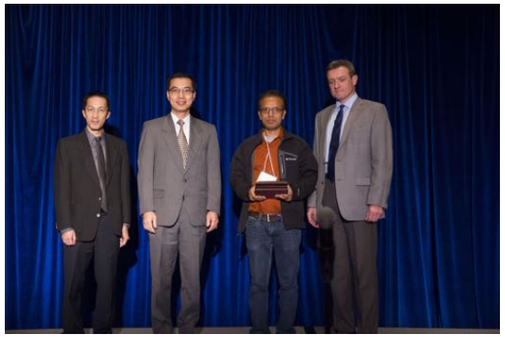

→ Click to read more and share

Best Student Paper

→ Click to read more and share

Give Feedback

Best Technology Showcase

→ Click to read more and share

Most Influential Paper

→ Click to read more and share

People's Choice Technology Showcase

→ Click to read more and share

Give Feedback

Pictures

Think: Cognitive Computing, Big Data, Cloud, Security and Privacy

October 29 - 31

CASCON News

Download the CASCON 2018 Program Guide

See the CASCON 2018 Proceeding in ACM

Get certain Workshops materials

Call for Papers for CASCON 2019

Give Feedback

Share the page:

Whole Team Picture

Five Years Team Picture

Ten Years Team Picture

→ Click to read more and share

→ Click to read more and share

→ Click to read more and share

Fifteen Years Team Picture

→ Click to read more and share

Twenty Years Team Picture

→ Click to read more and share

Twenty Five Years Team Picture

→ Click to read more and share

Give Feedback

Twenty Eight Years Team Picture

→ Click to read more and share

Volunteer's Team Picture

→ Click to read more and share

STEM 4 Girls Team Picture

→ Click to read more and share

Give Feedback

Archives

Think: Cognitive Computing, Big Data, Cloud, Security and Privacy

October 29 - 31

CASCON News

Download the CASCON 2018 Program Guide

See the CASCON 2018 Proceeding in ACM

Get certain Workshops materials

Call for Papers for CASCON 2019

Give Feedback

CASCON Proceedings Archives

- Proceedings Collection
- CASCON 2018

— CASCON 2013 - CASCON 2017

- CASCON 2017

CASCON Website Archives

- PDF CASCON 2018
- PDF CASCON 2017

— CASCON 2012 - CASCON 2016

- PDF CASCON 2016

CASCON 2016

CASCON 2015

CASCON 2014

CASCON 2013

— CASCON 2008 - CASCON 2012

CASCON 2012

CASCON 2011

CASCON 2010

CASCON 2009

CASCON 2008

— CASCON 2003 - CASCON 2007

CASCON 2007

CASCON 2006

CASCON 2005

CASCON 2004

CASCON 2003

— CASCON 1998 - CASCON 2002

CASCON 2002

PDF CASCON 2015

PDF CASCON 2014

PDF CASCON 2012

— CASCON 1998 - CASCON 2011

PDF CASCON 2011

PDF CASCON 2010

PDF CASCON 2009

PDF CASCON 2008

PDF CASCON 1998 - 2007

Give Feedback

- CASCON 2001
- CASCON 2000
- CASCON 1999
- CASCON 1998

— CASCON 1994 - CASCON 1997

- CASCON 1997
- CASCON 1996
- CASCON 1995
- CASCON 1994

— CASCON 1991 - CASCON 1993

- CASCON 1993 Volume-2
- CASCON 1993 Volume-1
- CASCON 1992 Volume-2
- CASCON 1992 Volume-1
- CASCON 1991

Give Feedback

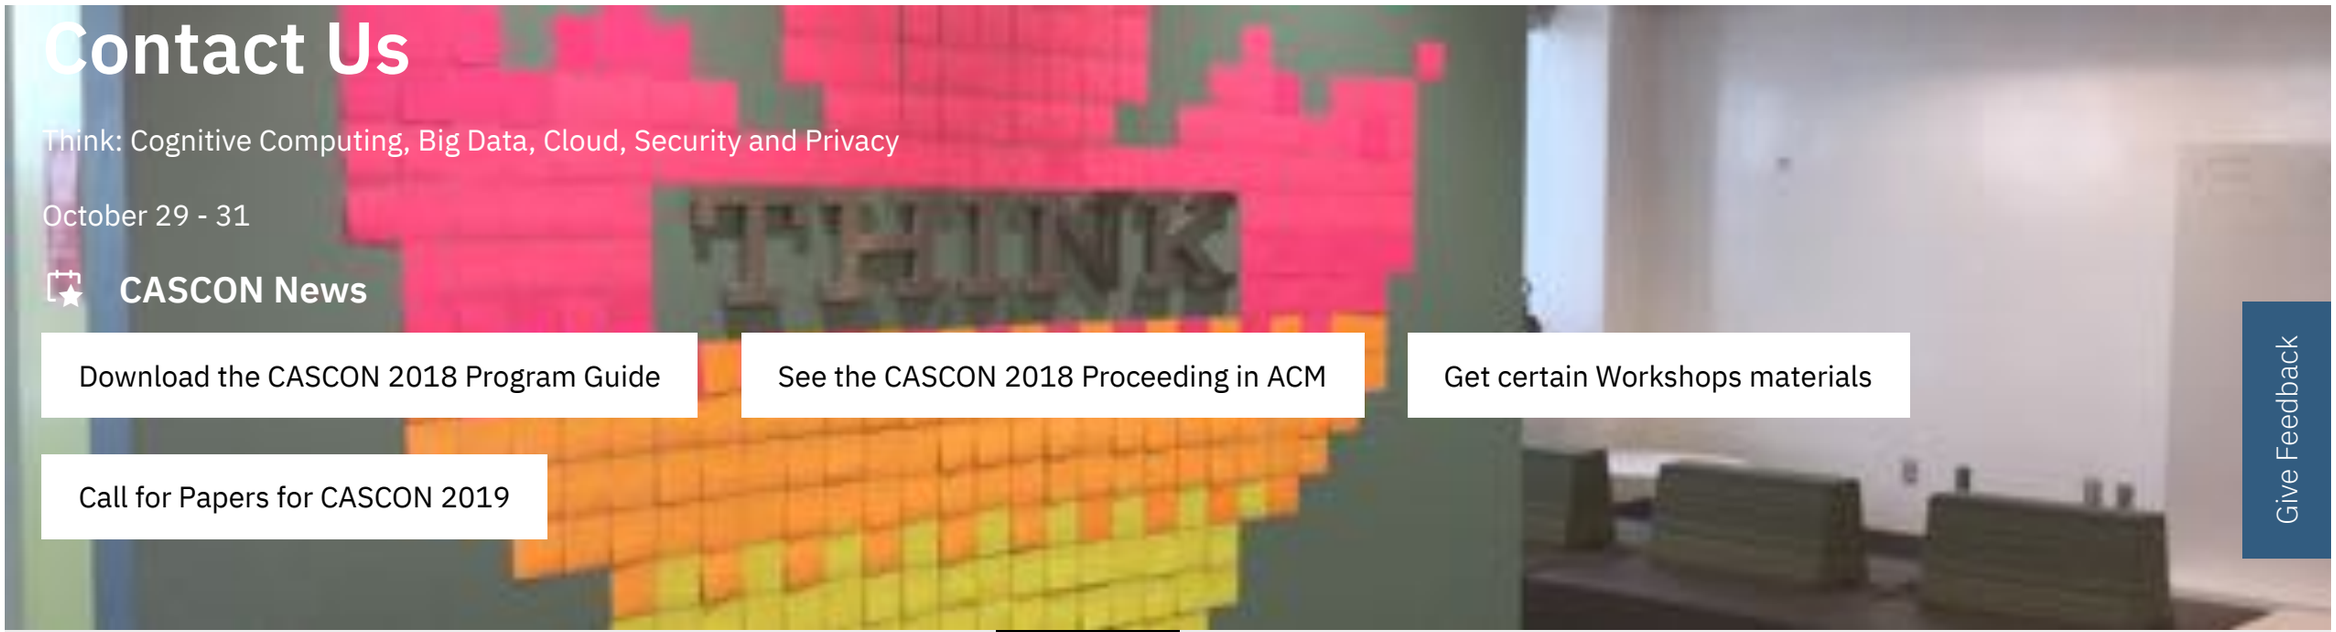

Questions, Suggestions, Concerns?

Email Us!

info@cascon.ca