NSERC Strategic Network

# "SMART APPLICATIONS ON VIRTUAL INFRASTRUCTURE"

## 2015 Annual General Meeting Agenda: Tuesday, July 7, 2015 - 8:25 am to 5:00 pm approx.

| | |
|---|---|
| **8:25 am** | *Welcome and AGM Overview* - **Al Leon-Garcia**, SAVI Scientific Director |
| **8:30 am** | *Big Data Analytics in SAVI* - **Marin Litoiu**, York University |
| **9:00 am** | *Challenges in Cloud Packet Capture* - **Nykolai Bilaniuk**, System Engineer, NIKSUN |
| **9:30 am** | *Hadoop Orchestration* - **Raouf Boutaba**, University of Waterloo |
| **10:00 am** | Poster and Demo Session 1 |
| **11:00 am** | *Keynote: NFV and Orchestration* - **Walter Miron**, Director of Technology Strategy, TELUS Communications Co. |
| **11:45 am** | *Technical Overview of SAVI SDI Testbed* - **Hadi Bannazadeh**, SAVI Chief Testbed Architect, University of Toronto |
| **12:15 pm** | Lunch - Poster and Demo Session 2 |
| **1:30 pm** | *Keynote: Smart City Challenges* - **Dragan Narandzic**, CTO, Ericsson Canada |
| **2:15 pm** | *From SAVI SDI to Smart City Platforms* - **Al Leon-Garcia**, SAVI Scientific Director |
| **2:45 pm** | Poster and Demo Session 3 |
| **3:30 pm** | *SDN-Enabled Service Automation* - **Robert Keys**, CTO, BTI Systems Inc. |
| **4:00 pm** | *Wireless over Fiber Access* - **Tho Le-Ngoc**, McGill University |
| **4:30 pm** | *CENI SDN Optical Testbed* - **Rod Wilson**, Senior Director External Research, Ciena |
| **5:00 pm** | *Closing Remarks* - **Al Leon-Garcia**, SAVI Scientific Director |

## SAVI Testbed Workshop Topics: Monday, July 6, 2015 - 9:00 am to 6:00 pm approx.

| | |
|---|---|
| **9:00 am** | Introduction to Basic Use of SAVI Testbed |
| **10:45 am** | Heat Orchestration: Autoscaling; Platform-wide |
| **12:15 pm** | Lunch |
| **1:00 pm** | NFV in SAVI SDI and ViNO (Virtual Network Overlay) |
| **2:45 pm** | SAVI/GENI Federation: Joint GENI/SAVI Experiment |
| **4:30 pm** | Big Data using Sahara and Spark |
| **6:00 pm** | Conclusion |

**Location:** Bahen Building, Room 1180, University of Toronto Campus, 40 St. George Street (Ground Floor)
**Contact:** Vladi Cirillo, Network Manager, (416) 946-3881, v.cirillo@utoronto.ca *www.savinetwork.ca*

# POSTER & DEMONSTRATION TITLE ROSTER

## THEME 1 - FUTURE APPLICATIONS

### POSTERS

**"*Wide Area Big Data Processing on the SAVI Testbed*"** Zhiming Hu and **Baochun Li,** (University of Toronto)

**"*Bellini: Application Overlay as a Service*"** Shuhao Liu, Yinan Liu, Siyi Wu, **Baochun Li,** (University of Toronto)

**"*Consistent Wide-Area Network Topology Updates*"** Yinan Liu, Shuhao Liu, **Baochun Li,** (University of Toronto)

**"*Beehive: Repairing Multiple Failures with Optimal Network Transfer over SAVI*"** Jun Li and **Baochun Li,** (University of Toronto)

**"*A Resource Allocation Mechanism for Video Mixing as a Cloud Computing Service in Multimedia Conferencing Applications*"** Abbas Soltanian and **Roch Glitho,** (Université Concordia University)

**"*A Platform Architecture for Multimedia Conferencing Applications in the Cloud*"** Ahmad Ferdous Bin Alam, Sami Yangui, **Roch Glitho,** (Université Concordia University)

**"*A Use Case on Network Function Virtualization for the Rapid Provisioning of new Video Delivery Services in Content Delivery Networks*"** Narjes Tahghigh, Sandhya Shanmugasundaram, Jagruti Sahoo, Sami Yangui, **Roch Glitho,** (Université Concordia University)

**"*From Relations to Multi-Dimensional Maps: Towards A SQL-to-HBase Transformation Methodology*"** Diego Serrano, Dan Han, **Eleni Stroulia**, (University of Alberta)

**"*A Cloud-Based Architecture for Real-Time Multi-modal Interactions and Analytics (Towards the Kaleidoscope App)*"** Hu Zhang, Diego Serrano, **Eleni Stroulia,** (University of Alberta)

### POSTER & DEMONSTRATION

**"*Multi-Party Video Conferencing with Bellini on the SAVI Testbed*"** Shuhao Liu, Yinan Liu, **Baochun Li**, (University of Toronto)

**"*Live Video Broadcast with Bellini on the SAVI Testbed*"** Shuhao Liu, Yinan Liu, **Baochun Li,** (University of Toronto)

**"*Kaleidoscope --- Applications for Multi-Tier Cloud, Smart Edge/Core and Adaptive Resource Support*"** Ronald Desmarais, Andreas Bergen**, Hausi Müller, Sudhakar Ganti,** (University of Victoria)

# POSTER & DEMONSTRATION TITLE ROSTER

### DEMONSTRATIONS

**"*The Ignite Distributed Collaborative Scientific Visualization System*"** Sushil Bhojwani, Matt Hemmings, Rick McGeer, Yvonne Coady, Ulrike Stege, (University of Victoria), Glenn Ricart, (US Ignite), David Lary, (University of Texas at Dallas), Jens Lincke, (Hasso-Plattner Institute), Dan Ingalls, Robert Krahn, Marko Roeder, (Communication Design Group, SAP)

## THEME 2 - RESOURCE CONTROL & MANAGEMENT

### POSTERS

**"*EdgeX: Edge Replication for Web Applications*"** Hemant Saxena, **Ken Salem**, (University of Waterloo)

**"*Self Partitioning, Edge-Aware Replicated DataBase*"** Cătălin Avram and **Ken Salem**, (University of Waterloo)

**"*Lightweight Robust Optimizer for Distributed Application Deployment in Multi-Clouds*"** Ravneet Kaur, **Murray Woodside, John Chinneck,** (Carleton University)

**"*Evaluating Cluster Configurations for Big Data Processing: An Exploratory Study*"** Roni Sandel, Marios Fokaefs, Mark Shtern, **Marin Litoiu,** (York University)

**"*Engineering Economics of Everything-as-a-Service (XaaS)*"** Marios Fokaefs and **Marin Litoiu,** (York University)

### POSTER & DEMONSTRATION

**"*OpenFlow-enabled DDoS Attack Mitigation on Hybrid Cloud*"** Nasim Beigi-Mohammadi, Cornel Barna, Hamzeh Khazei, **Marin Litoiu,** (York University)

**"*K-Feed, A Data-Oriented Approach to Application Performance Management in Cloud*"** Saeed Zareian, Rodrigo Veleda, Hamzeh Khazei, **Marin Litoiu,** (York University)

### DEMONSTRATIONS

**"*Apache Spark on SAVI with OpenStack Sahara*"** Hamzeh Khazaei, Marios Fokaefs, **Marin Litoiu,** (York University)

**"*Enabling an Enhanced Data-as-a-Service Ecosystem*"** Justin Cuaresma, Ethan Nguyen, Mark Shtern, **Marin Litoiu,** (York University)

**"*Hogna: A Platform for Developing Self-Adaptive Applications in Cloud*"** Cornel Barna and **Marin Litoiu,** (York University)

# POSTER & DEMONSTRATION TITLE ROSTER

## THEME 3 - SMART CONVERGED EDGE

### POSTERS

**"*Robust Resource Reservation in Virtual Wireless Networks*"** Ali Abbasi and **Majid Ghaderi,** (University of Calgary)

**"*Simple Distributed Programming in Software-Defined Networks*"** Soheil Hassas Yeganeh and **Yashar Ganjali,** (University of Toronto)

**"*Dynamic Proactive Routing in Software Defined Networks*"** Sajad Shirali-Shahreza and **Yashar Ganjali,** (University of Toronto)

**"*Dynamic Resource Allocation for MapReduce with Partitioning Skew*"** Zhihong Liu (Univ. of Waterloo), Qi Zhang (Univ. of Waterloo & Toronto), **Raouf Boutaba** (Univ. of Waterloo), Yaping Liu and Zhenghu Gong (National University of Defense Technology)

**"*SiMPLE: Survivability in MultiPath*"** Md Mashrur Alam Khan, Nashid Shahriar, Reaz Ahmed, **Raouf Boutaba,** (University of Waterloo)

**"*nf.io: A File System Abstraction for NFV Orchestration*"** Md. Faizul Bari, Shihabur Rahman Chowdhury, Reaz Ahmed, **Raouf Boutaba,** (University of Waterloo)

**"*Aurora: Adaptive Block Replication in Distributed File Systems*"** Qi Zhang (Univ. of Waterloo & Toronto), Sai Qian Zhang, **Alberto Leon-Garcia,** (University of Toronto), **Raouf Boutaba**, (University of Waterloo)

**"*Network Function Virtualization Enabled Multicast Routing on SDN*"** Sai Qian Zhang, Qi Zhang, Hadi Bannazadeh, **Alberto Leon-Garcia** (University of Toronto)

**"*Fast Network Flow Resumption for Live Virtual Machine Migration on SDN*"** Sai Qian Zhang, Pouya Yasrebi, Ali Tizghadam, Hadi Bannazadeh, **Alberto Leon-Garcia,** (University of Toronto)

"*Anomaly Detection and Localization in Software Defined Infrastructures*" Qi Zhang, Jieyu (Eric) Lin, Hadi Bannazadeh, **Alberto Leon-Garcia,** (University of Toronto)

### POSTER & DEMONSTRATION

**"*Multicast as a Service On SAVI Network*"** Sai Qian Zhang, Hadi Bannazadeh, **Alberto Leon-Garcia,** (University of Toronto)

# POSTER & DEMONSTRATION TITLE ROSTER

## THEME 4 - INTEGRATED WIRELESS/OPTICAL ACCESS

### POSTERS

*"Leveraging Synergy of SDWN and Multi-Layer Resource Management for Network Optimization"* Mahsa Derakhshani, Saeedeh Parsaeefard, **Tho Le-Ngoc,** (McGill University), **Alberto Leon-Garcia,** (University of Toronto)

*"Virtualization of Multi-Cell 802.11 Networks: Association and Airtime Control"* Mahsa Derakhshani, Xiaowei Wang, **Tho Le-Ngoc,** (McGill University), **Alberto Leon-Garcia,** (University of Toronto)

*"Adaptive Pilot Duration Allocation in Wireless Virtualized Networks with Massive MIMO"* Rajesh Dawadi, Saeedeh Parsaeefard, Mahsa Derakhshani, **Tho Le-Ngoc,** (McGill University)

*"Full Duplex Analog WiFi over PONs"* Zhihui Cao, An T. Nguyen, **Leslie A. Rusch,** (Université Laval)

*"Virtualized WiFi on Radio Over Fiber for PONs"* Zhihui Cao, Robert Morawski, Quang Dong Ho, Thanh Ngon Tran, **Tho Le-Ngoc**, (McGill University), **Leslie A. Rusch,** (Université Laval)

### POSTER & DEMONSTRATION

*"Virtualized and Software-Defined WiFi Radio Access Networks"* Quang-Dung Ho, Alfred Kenny, **Tho Le-Ngoc,** (McGill University)

## THEME 5 - APPLICATION PLATFORM TESTBED

### POSTERS

*"Expanding OpenFlow Capabilities with Virtualized Reconfigurable Hardware"* Stuart Byma, Naif Tarafdar, Talia Xu, Hadi Bannazadeh, **Alberto Leon-Garcia, Paul Chow,** (University of Toronto)

*"OpenCLoud: A Heterogenous Reconfigurable Cloud Platform with an OpenCL Front-end"* Naif Tarafdar, Eric Fukuda, William Suriaputra, **Paul Chow**, (University of Toronto)

*"Interactive Large-scale Data Analysis on Virtual FPGAs in the Cloud"* Eric Fukuda, Naif Tarafdar, **Paul Chow,** (University of Toronto)

*"SDI Manager Architecture"* Byungchul Park, Thomas Lin, Hadi Bannazadeh, **Alberto Leon-Garcia,** (University of Toronto)

*"Enhancing VoIP Data Plane using SDN"* Lilin Zhang, Ali Tizghadam, Hadi Bannazadeh, **Alberto Leon-Garcia,** (University of Toronto)

*"Operations, Maintenance and Diagnostics for SAVI Testbed SDI"* Joseph Wahba, Hadi Bannazadeh, **Alberto Leon-Garcia,** (University of Toronto)

# POSTER & DEMONSTRATION TITLE ROSTER

## POSTERS: Theme 5 Cont'd

***"Idle VM Consolidation in SAVI Testbed"*** SeyedAli Jokar, Joseph Wahba, Hadi Bannazadeh, **Alberto Leon-Garcia,** (University of Toronto)

***"Security Function Virtualization in Software Defined Infrastructure"*** Pouya Yasrebi, Sina Monfared, Hadi Bannazadeh, **Alberto Leon-Garcia,** (University of Toronto)

***"Scalability in Security Function Virtualization of Software Defined Infrastructure"*** Pouya Yasrebi, Sai Zhang, Hadi Bannazadeh, **Alberto Leon-Garcia,** (University of Toronto)

***"SAVI-GENI Cloud Federation: Resource Procurement"*** Sushil Bhojwani, Riz Panjwani, Rick McGeer, **Sudhakar Ganti,** Andi Bergen, Ron Desmarais, **Hausi Müller,** (University of Victoria), Hadi Bannazadeh (University of Toronto)

***"SAVI-GENI Cloud Federation: Authentication"*** Sushil Bhojwani, Riz Panjwani, Rick McGeer, **Sudhakar Ganti,** Andi Bergen, Ron Desmarais, **Hausi Müller,** (University of Victoria)

## POSTER & DEMONSTRATION

***"MonArch Multilayer Monitoring and Analytics in Software-Defined Infrastructures"*** Jieyu (Eric) Lin, Hadi Bannazadeh, **Alberto Leon-Garcia,** (University of Toronto)

***"Software Defined Infrastructure overlay over heterogeneous cloud to facilitate NFV service chaining"*** Spandan Bemby, SeyedAli Jokar, Hadi Bannazadeh, **Alberto Leon-Garcia,** (University of Toronto)

***"Flexible NFV orchestration using Containerized Overlay"*** SeyedAli Jokar, Spandan Bemby, Hadi Bannazadeh, **Alberto Leon-Garcia,** (University of Toronto)

***"NFV Realization on Virtual Customer Premise Edge (vCPE) using SAVI SDI"*** Sina Monfared, Hadi Bannazadeh, **Alberto Leon-Garcia,** (University of Toronto)

## DEMONSTRATION

***"SAVI-GENI Cloud Federation"*** Sushil Bhojwani, Riz Panjwani, Rick McGeer, **Sudhakar Ganti,** Andi Bergen, Ron Desmarais, **Hausi Müller,** (University of Victoria), Hadi Bannazadeh, (University of Toronto)

# POSTER & DEMONSTRATION SCHEDULE

## 10:00 A.M. - SESSION 1: POSTERS

| STN.NO. | POSTER TITLE |
|---|---|
| P1 | *"EdgeX: Edge Replication for Web Applications"* Hemant Saxena, **Ken Salem,** (University of Waterloo) |
| P2 | *"Self Partitioning, Edge-Aware Replicated DataBase"* Cătălin Avram and **Ken Salem,** (University of Waterloo) |
| P3 | *"Lightweight Robust Optimizer for Distributed Application Deployment in Multi-Clouds"* Ravneet Kaur, **Murray Woodside, John Chinneck,** (Carleton University) |
| P4 | *"Evaluating Cluster Configurations for Big Data Processing: An Exploratory Study"* Roni Sandel, Marios Fokaefs, Mark Shtern, **Marin Litoiu,** (York University) |
| P5 | *"Engineering Economics of Everything-as-a-Service (XaaS)"* Marios Fokaefs and **Marin Litoiu,** (York University) |
| P6 | *"Robust Resource Reservation in Virtual Wireless Networks"* Ali Abbasi and **Majid Ghaderi,** (University of Calgary) |
| P7 | *"Simple Distributed Programming in Software-Defined Networks"* Soheil Hassas Yeganeh and **Yashar Ganjali,** (University of Toronto) |
| P8 | *"Dynamic Proactive Routing in Software Defined Networks"* Sajad Shirali-Shahreza and **Yashar Ganjali,** (University of Toronto) |
| P9 | *"Dynamic Resource Allocation for MapReduce with Partitioning Skew"* Zhihong Liu (Univ. of Waterloo), Qi Zhang (Univ. of Waterloo & Toronto), **Raouf Boutaba,** (Univ. of Waterloo), Yaping Liu and Zhenghu Gong (National University of Defense Technology) |
| P10 | *"SiMPLE: Survivability in MultiPath"* Md Mashrur Alam Khan, Nashid Shahriar, Reaz Ahmed, **Raouf Boutaba,** (University of Waterloo) |
| P11 | *"nf.io: A File System Abstraction for NFV Orchestration"* Md. Faizul Bari, Shihabur Rahman Chowdhury, Reaz Ahmed, **Raouf Boutaba,** (University of Waterloo) |
| P12 | *"Aurora: Adaptive Block Replication in Distributed File Systems"* Qi Zhang (Univ. of Waterloo & Toronto), Sai Qian Zhang, **Alberto Leon-Garcia,** (University of Toronto), **Raouf Boutaba,** (University of Waterloo) |
| P13 | *"Network Function Virtualization Enabled Multicast Routing on SDN"* Sai Qian Zhang, Qi Zhang, Hadi Bannazadeh, **Alberto Leon-Garcia,** (University of Toronto) |
| P14 | *"Fast Network Flow Resumption for Live Virtual Machine Migration on SDN"* Sai Qian Zhang, Pouya Yasrebi, Ali Tizghadam, Hadi Bannazadeh, **Alberto Leon-Garcia,** (University of Toronto) |

# POSTER & DEMONSTRATION SCHEDULE

| STN.NO. | POSTER TITLE: 10:00 am Session 1 Cont'd |
|---|---|
| P15 | *"Anomaly Detection and Localization in Software Defined Infrastructures"* Qi Zhang, Jieyu (Eric) Lin, Hadi Bannazadeh, **Alberto Leon-Garcia,** (University of Toronto) |

## 10:00 A.M. - SESSION 1: POSTERS & DEMONSTRATIONS

| STN.NO. | POSTER & DEMONSTRATION TITLE |
|---|---|
| PD1 | *"OpenFlow-enabled DDoS Attack Mitigation on Hybrid Cloud"* Nasim Beigi-Mohammadi, Cornel Barna, Hamzeh Khazei, **Marin Litoiu,** (York University) |
| PD2 | *"K-Feed, A Data-Oriented Approach to Application Performance Management in Cloud"* Saeed Zareian, Rodrigo Veleda, Hamzeh Khazei, **Marin Litoiu,** (York University) |
| PD3 | *"Multicast as a Service On SAVI Network"* Sai Qian Zhang, Hadi Bannazadeh, **Alberto Leon-Garcia,** (University of Toronto) |

## 10:00 A.M. - SESSION 1: DEMONSTRATIONS

| STN.NO. | DEMONSTRATION TITLE |
|---|---|
| D1 | *"Apache Spark on SAVI with OpenStack Sahara"* Hamzeh Khazaei, Marios Fokaefs, **Marin Litoiu,** (York University) |
| D2 | *"Enabling an Enhanced Data-as-a-Service Ecosystem"* Justin Cuaresma, Ethan Nguyen, Mark Shtern, **Marin Litoiu,** (York University) |
| D3 | *"Hogna: A Platform for Developing Self-Adaptive Applications in Cloud"* Cornel Barna and **Marin Litoiu,** (York University) |

## 12:00 NOON - SESSION 2: POSTERS

| STN.NO. | POSTER TITLE |
|---|---|
| P1 | *"Leveraging Synergy of SDWN and Multi-Layer Resource Management for Network Optimization"* Mahsa Derakhshani, Saeedeh Parsaeefard, **Tho Le-Ngoc,** (McGill University), **Alberto Leon-Garcia,** (University of Toronto) |
| P2 | *"Virtualization of Multi-Cell 802.11 Networks: Association and Airtime Control"* Mahsa Derakhshani, Xiaowei Wang, **Tho Le-Ngoc,** (McGill University), **Alberto Leon-Garcia,** (University of Toronto) |

# POSTER & DEMONSTRATION SCHEDULE

| STN.NO. | POSTER TITLE: Noon Session 2 Cont'd |
|---|---|
| P3 | *"Adaptive Pilot Duration Allocation in Wireless Virtualized Networks with Massive MIMO"* Rajesh Dawadi, Saeedeh Parsaeefard, Mahsa Derakhshani, **Tho Le-Ngoc,** (McGill University) |
| P4 | *"Full Duplex Analog WiFi over PONs"* Zhihui Cao, An T. Nguyen, **Leslie A. Rusch,** (Université Laval) |
| P5 | *"Virtualized WiFi on Radio Over Fiber for PONs"* Zhihui Cao, Robert Morawski, Quang Dong Ho, Thanh Ngon Tran, **Tho Le-Ngoc,** (McGill University), **Leslie A. Rusch,** (Université Laval) |
| P6 | *"Expanding OpenFlow Capabilities with Virtualized Reconfigurable Hardware"* Stuart Byma, Naif Tarafdar, Talia Xu, Hadi Bannazadeh, **Alberto Leon-Garcia, Paul Chow,** (University of Toronto) |
| P7 | *"OpenCLoud: A Heterogenous Reconfigurable Cloud Platform with an OpenCL Front-end"* Naif Tarafdar, Eric Fukuda, William Suriaputra, **Paul Chow,** (University of Toronto) |
| P9 | *"Interactive Large-scale Data Analysis on Virtual FPGAs in the Cloud"* Eric Fukuda, Naif Tarafdar, **Paul Chow,** (University of Toronto) |
| P10 | *"SDI Manager Architecture"* Byungchul Park, Thomas Lin, Hadi Bannazadeh, **Alberto Leon-Garcia,** (University of Toronto) |
| P11 | *"Enhancing VoIP Data Plane using SDN"* Lilin Zhang, Ali Tizghadam, Hadi Bannazadeh, **Alberto Leon-Garcia,** (University of Toronto) |
| P12 | *"Operations, Maintenance and Diagnostics for SAVI Testbed SDI"* Joseph Wahba, Hadi Bannazadeh, **Alberto Leon-Garcia,** (University of Toronto) |
| P13 | *"Idle VM Consolidation in SAVI Testbed"* SeyedAli Jokar, Joseph Wahba, Hadi Bannazadeh, **Alberto Leon-Garcia,** (University of Toronto) |
| P14 | *"Security Function Virtualization in Software Defined Infrastructure"* Pouya Yasrebi, Sina Monfared, Hadi Bannazadeh, **Alberto Leon-Garcia,** (University of Toronto) |
| P15 | *"Scalability in Security Function Virtualization of Software Defined Infrastructure"* Pouya Yasrebi, Sai Zhang, Hadi Bannazadeh, **Alberto Leon-Garcia,** (University of Toronto) |

## 12:00 NOON - SESSION 2: POSTER & DEMONSTRATION

| STN.NO. | POSTER & DEMONSTRATION TITLE |
|---|---|
| PD1 | *"Virtualized and Software-Defined WiFi Radio Access Networks"* Quang-Dung Ho, Alfred Kenny, **Tho Le-Ngoc,** (McGill University) |
| PD2 | *"MonArch Multilayer Monitoring and Analytics in Software-Defined Infrastructures"* Jieyu (Eric) Lin, Hadi Bannazadeh, **Alberto Leon-Garcia,** (University of Toronto) |

# POSTER & DEMONSTRATION SCHEDULE

| STN.NO. | POSTER & DEMONSTRATION TITLE : Noon Session 2 Cont'd |
|---|---|
| PD3 | *"Software Defined Infrastructure overlay over heterogeneous cloud to facilitate NFV service chaining"* Spandan Bemby, SeyedAli Jokar, Hadi Bannazadeh, **Alberto Leon-Garcia,** (University of Toronto) |
| PD4 | *"Flexible NFV orchestration using Containerized Overlay"* SeyedAli Jokar, Spandan Bemby, Hadi Bannazadeh, **Alberto Leon-Garcia,** (University of Toronto) |
| PD5 | *"NFV Realization on Virtual Customer Premise Edge (vCPE) using SAVI SDI"* Sina Monfared, Hadi Bannazadeh, **Alberto Leon-Garcia,** (University of Toronto) |

## 2:45 P.M. - SESSION 3: POSTERS

| STN.NO. | POSTER TITLE |
|---|---|
| P1 | "*Wide Area Big Data Processing on the SAVI Testbed*" Zhiming Hu and **Baochun Li,** (University of Toronto) |
| P2 | "*Bellini: Application Overlay as a Service*" Shuhao Liu, Yinan Liu, Siyi Wu, **Baochun Li,** (University of Toronto) |
| P3 | "*Consistent Wide-Area Network Topology Updates*" Yinan Liu, Shuhao Liu, **Baochun Li,** (University of Toronto) |
| P4 | "*Beehive: Repairing Multiple Failures with Optimal Network Transfer over SAVI*" Jun Li and **Baochun Li,** (University of Toronto) |
| P5 | "*A Resource Allocation Mechanism for Video Mixing as a Cloud Computing Service in Multimedia Conferencing Applications*" Abbas Soltanian and **Roch Glitho,** (Université Concordia University) |
| P6 | "*A Platform Architecture for Multimedia Conferencing Applications in the Cloud*" Ahmad Ferdous Bin Alam, Sami Yangui, **Roch Glitho,** (Université Concordia University) |
| P7 | "*A Use Case on Network Function Virtualization for the Rapid Provisioning of new Video Delivery Services in Content Delivery Networks*" Narjes Tahghigh, Sandhya Shanmugasundaram, Jagruti Sahoo, Sami Yangui, **Roch Glitho,** (Université Concordia University) |
| P8 | "*From Relations to Multi-Dimensional Maps: Towards A SQL-to-HBase Transformation Methodology*" Diego Serrano, Dan Han, **Eleni Stroulia,** (University of Alberta) |
| P9 | "*A Cloud-Based Architecture for Real-Time Multi-modal Interactions and Analytics (Towards the Kaleidoscope App)*" Hu Zhang, Diego Serrano, **Eleni Stroulia,** (University of Alberta) |

# POSTER & DEMONSTRATION SCHEDULE

## 2:45 P.M. - SESSION 3: POSTER & DEMONSTRATION

| STN.NO. | POSTER & DEMONSTRATION TITLE |
|---|---|
| PD1 | **"*Multi-Party Video Conferencing with Bellini on the SAVI Testbed*"** Shuhao Liu, Yinan Liu, **Baochun Li,** (University of Toronto) |
| PD2 | **"*Live Video Broadcast with Bellini on the SAVI Testbed*"** Shuhao Liu, Yinan Liu, **Baochun Li,** (University of Toronto) |
| PD3 | **"*Kaleidoscope --- Applications for Multi-Tier Cloud, Smart Edge/Core and Adaptive Resource Support*"** Ronald Desmarais, Andreas Bergen, **Hausi Müller, Sudhakar Ganti,** (University of Toronto) |

## 2:45 P.M. - SESSION 3: DEMONSTRATIONS

| STN.NO. | DEMONSTRATION TITLE |
|---|---|
| D1 | ***"The Ignite Distributed Collaborative Scientific Visualization System"*** Sushil Bhojwani, Matt Hemmings, Rick McGeer, Yvonne Coady, Ulrike Stege, (University of Victoria), Glenn Ricart, (US Ignite), David Lary, (University of Texas at Dallas), Jens Lincke, (Hasso-Plattner Institute), Dan Ingalls, Robert Krahn, Marko Roeder, (Communication Design Group, SAP) |
| D2 | ***"SAVI-GENI Cloud Federation"*** Sushil Bhojwani, Riz Panjwani, Rick McGeer, **Sudhakar Ganti,** Andi Bergen, Ron Desmarais, **Hausi Müller,** (University of Victoria), Hadi Bannazadeh (University of Toronto) |
| See Note | *"SAVI-GENI Cloud Federation: Resource Procurement"* Sushil Bhojwani, Riz Panjwani, Rick McGeer, **Sudhakar Ganti,** Andi Bergen, Ron Desmarais, **Hausi Müller,** (University of Victoria), Hadi Bannazadeh, (University of Toronto) |
| See Note | *"SAVI-GENI Cloud Federation: Authentication"* Sushil Bhojwani, Riz Panjwani, Rick McGeer, **Sudhakar Ganti,** Andi Bergen, Ron Desmarais, **Hausi Müller**, (University of Victoria) |

**Note:** The demonstration ***"SAVI-GENI Cloud Federation"*** is a combined working of the posters entitled ***"SAVI-GENI Cloud Federation: Resource Procurement"*** and ***"SAVI-GENI Cloud Federation: Authentication"***. The posters are separate works that will be presented in conjunction with the demonstration.

# GUEST SPEAKER BIO

**Mr. Nykolai Bilaniuk, PhD, P.Eng**
*System Engineer,*
NIKSUN Inc.

Nykolai started his career in research but later turned to more applied work in the private sector. His research experience has included acoustic sensing and optical communications at the National Research Council of Canada's Institute for Microstructural Sciences, and integrated circuits for mobile communications air interfaces at the Department of Electronics at Carleton University. Next, he spent time at Altera Corporation working on telecommunications intellectual property cores for FPGAs. For about a decade starting in 2001, he was with Comgate Engineering Ltd working on Canadian federal government and international projects dealing with communications security. Since 2012 he has been a system engineer with NIKSUN, a provider of network packet capture and analytics systems.

Nykolai Bilaniuk received his BASc in Engineering Science and MASc in Electrical Engineering from the University of Toronto in 1985 and 1988 respectively, and a PhD in Electrical and Computer Engineering from Carnegie Mellon University in 1992.

# GUEST SPEAKER BIO

**Mr. Walter Miron**
*Director of Technology Strategy,*
TELUS Communications Co.

Walter Miron is a Director or Technology Strategy at TELUS Communications, where he is responsible for the evolution of their TCP/IP and Optical networks.

He has over 20 years of experience in enterprise and service provider networking conducting technology selection and service development projects. Walter is a member of the research program committee of the SAVI project, the Heavy Reading Global Ethernet Executive Council, the ATOPs SDN/NFV working group, and represents TELUS at the Venus Cybersecurity Corporation, as the chairman and a board member for the Center of Excellence for Next Generation Networking (CENGN). He is frequently a speaker at industry conferences and working groups.

# GUEST SPEAKER BIO

**Dr. Hadi Bannazadeh**
*SAVI Chief Testbed Architect,*
University of Toronto

Dr. Bannazadeh received his MASc and BASc from Amir Kabir University of Technology (Tehran Polytechnic), Tehran, Iran.

Hadi holds a PhD from the University of Toronto's Department of Electrical & Computer Engineering. After graduating, he worked at Cisco Systems as a Senior Network Software Engineer. In 2011, he returned to the University of Toronto to lead the efforts towards the creation of Canadian national testbed as part of the Smart Applications on Virtual Infrastructure (SAVI) research project. Since then, he has been the Chief Testbed Architect for the SAVI project. Hadi's main research interest is in the field of Software Defined Infrastructure (SDI) including Software Defined Networking (SDN) and Cloud Computing.

Prior to beginning the University of Toronto's PhD program, Hadi was first employed as a Software Engineer and then as a Software Architect for the Iran Telecom Research Center and Iran Communications Industries Inc. where he contributed to the design and development of a large-scale telecomm switching system.

**List of Publications:**
http://dblp.uni-trier.de/pers/hd/b/Bannazadeh:Hadi

# GUEST SPEAKER BIO

**Mr. Dragan Nerandzic**
*Chief Technology Officer,*
Ericsson Canada

Dragan Nerandzic joined Ericsson in 2001 and since 2006 has held the role of Chief Technology Officer for Ericsson in Canada. Mr. Nerandzic previously held the roles of VP Network Systems and Director of Technical Strategy.

Prior to joining Ericsson, Mr. Nerandzic was responsible for wireless technology planning and strategy for an operator's network. He actively participated in standardization and industry organizations including 3GPP2, TIA and CDG. Mr. Nerandzic also had responsibilities for engineering design of analog and digital mobile networks. Mr. Nerandzic's past experience includes working at the Faculty of Electrical Engineering at the University of Belgrade where he held various teaching and research roles in areas of fixed and mobile networks.

Mr. Nerandzic is a member of the MET Advisory Board at the University of Toronto, the Editorial Advisory Board of Backbone Magazine and Member of the Technical Panel at the Terrestrial Wireless Systems Branch of Communication Research Centre Canada. He is also a member of the Board of Directors of the Mobile Experience Innovation Center.

Mr. Nerandzic holds a Master of Science degree in Telecommunications and Computer Networks, Protocols and Switching, and a Bachelor of Science degree in Electrical Engineering from the University of Belgrade.

For more information, contact:
Patricia MacLean
Director, Communications & Marketing
Ericsson Canada Inc.
Phone +1 (905) 206-7928
Mobile +1 (416) 414-7755
patricia.maclean@ericsson.com

# GUEST SPEAKER BIO

**Mr. Robert Keys,**
*Chief Technology Officer,*
BTI Systems Inc.

Dr. Robert Keys brings over 20 years of communications engineering and leadership experience to his role as Chief Technology Officer. Prior to joining BTI Systems, Robert served as Chief Engineer for Bookham (now Oclaro), defining the company's technical roadmap and strategic investment requirements. Prior to Bookham, Robert held various influential roles in the development organization at Nortel, where he was responsible for the successful development and delivery of multiple optical products to market.

Robert is a well-known authority in the optical industry, and is a frequently sought-after speaker. He has more than 12 patents granted in the area of optical communications.

Robert has a Ph.D & M.Sc in Electronics and Electrical Engineering and a B.Sc in Physics from the University of Glasgow.

# GUEST SPEAKER BIO

**Mr. Rodney Wilson,**
*Senior Director, External Research*
Ciena Canada

Mr. Wilson is responsible for Ciena's leadership & interactions with Universities and the research community, including national research and education networks. Residing within Ciena's CTO organization, he is able to orchestrate intersections between emerging technologies and research network experiments. He is a frequent contributor to research projects, demonstrations and discussions about advanced optical telecommunications systems.

Prior to his current role, Mr. Wilson was a senior advisor for the CTO at Nortel, and held other advanced technology roles during 13 years with the company, including director of Broadband Switching, and optical Ethernet development.

Mr. Wilson was originally trained in Electrical Engineering at Ryerson University in Toronto, Ontario and the University of Toronto. He is also a graduate of the Executive Management School at Stanford University in Palo Alto, California.

# ABSTRACTS FOR POSTERS & DEMONSTRATIONS

## THEME 1 - FUTURE APPLICATIONS

### ➢ POSTERS:

- ***Wide Area Big Data Processing on the SAVI Testbed***

Zhiming Hu and **Baochun Li**, (University of Toronto)

Big data processing are widely used to distill insights from large volumes of data using learning algorithms. As both data and datacenters are geographically distributed, traditional solutions need to gather all their data across geographically distributed datacenters and analyze them in a centralized manner. Since the bandwidth across datacenters is expensive and data size is growing data grows exponentially, it is no longer efficient, nor practical to do this. In this work, we propose that data should be analyzed directly across geographically distributed datacenters, and demonstrate the advantages of this approach. We present our measurement studies and a new optimization framework to illustrate the feasibility of processing big data on the SAVI testbed, which is a geographically distributed computation facility and is thus an ideal candidate for wide area big data processing.

- ***Bellini: Application Overlay as a Service Networks***

Shuhao Liu, Yinan Liu, Siyi Wu, **Baochun Li**, (University of Toronto)

The SAVI testbed regions are interconnected by dedicated links or tunnels, which offer low-latency and high-bandwidth data transfer for inter-region communication. These are critical resources for the users. In order to fully utilize these resources, we propose Bellini, an efficient, highly flexible, and customizable framework, providing a data delivery service including a software-defined application overlay across the SAVI regions.

Bellini takes the back-end responsibilities of using an application-layer software-defined infrastructure to adapt to resource availability on the fly during runtime. It makes it feasible for smart applications to achieve high utilization of multiple types of resources on the SAVI testbed with ease. In addition, the service provided by Bellini is completely transparent. Smart applications can easily integrate with this service by delivering their data to a Bellini server via standard protocols, e.g., HTTP and RTMP. The traffic is thus delivered via the application overlay, and end users do not need to be concerned with any further details of implementing the service.

- ***Consistent Wide-Area Network Topology Updates***

Yinan Liu, Shuhao Liu, **Baochun Li**, (University of Toronto)

The SAVI testbed is a wide-area infrastructure that allows data to be transmitted and processed across geographically distributed datacenters. Since it is accessible by any users at any time, the traffic demands vary dynamically. It is therefore imperative to update the application-layer network topology to maintain adequate performance, such as meeting traffic deadlines and reducing packet losses. However, such configuration changes may easily lead to critical failures. The emergence of the software-defined networking paradigm provides a logically centralized controller, which helps us simplify the process of network updates in application-layer network topologies.

Despite the merits of using software-defined infrastructures, existing solutions are not able to solve two main problems satisfactorily: (1) finding an appropriate update sequence; and (2) rule space limitations on the network switch. In this work, we present a new technique that completes the wide area network topology updating at the application layer, solving both problems while keeping packet consistency during the update.

- ***Beehive: Repairing Multiple Failures with Optimal Network Transfer over SAVI***

Jun Li and **Baochun Li**, (University of Toronto)

Distributed storage systems have been increasingly deploying erasure codes (such as Reed-Solomon codes) for fault tolerance. Although Reed-Solomon codes require much less storage space than replication, a significant amount of network transfer and disk I/O will be imposed when fixing unavailable data by reconstruction. Traditionally, it is expected that unavailable data are fixed separately. However, since it is observed that failures can be correlated, fixing unavailable data of multiple failures is both unavoidable and even common. In this paper, we show that reconstructing data of multiple failures in batches can cost significantly less network transfer and disk I/O than fixing them separately. We present Beehive, a new design of erasure codes that can fix unavailable data of multiple failures in batches while consuming optimal network transfer with nearly optimal storage overhead. Evaluation results on the SAVI platform show that Beehive codes can save network transfer by up to 69.4% and disk I/O by 75% during reconstruction.

- ***A Resource Allocation Mechanism for Video Mixing as a Cloud Computing Service in Multimedia Conferencing Applications***

Abbas Soltanian and **Roch Glitho**, (Université Concordia University)

Multimedia conferencing is the conversational exchange of multimedia content between multiple parties. It has a wide range of applications [e.g. Massively Multiplayer Online Games (MMOGs) and distance learning]. Many multimedia conferencing applications use video extensively, thus video mixing in conferencing settings is of critical importance. This poster proposes a solution to optimize resource allocation for cloud-based video mixing service in multimedia conferencing applications, which can support scalability in terms of number of users, while guaranteeing Quality of Service (QoS). We formulate the resource allocation problem mathematically as an Integer Linear Programming (ILP) problem and design a heuristic for it. Simulation results show that our resource allocation model can support more participants compared to state-of-the-art, while honoring QoS, with respect to end-to-end delay.

- ***A Platform Architecture for Multimedia Conferencing Applications in the Cloud***

Ahmad Ferdous Bin Alam, Sami Yangui, **Roch Glitho**, (Université Concordia University)

Multimedia conferencing is the conversational exchange of multimedia content between multiple parties. It has a wide range of applications (e.g. Massively Multiplayer Online Games (MMOGs) and distance learning). In this poster, we propose an architecture for a Platform as-a-Service (PaaS) which (i) supports the whole provisioning process of conferencing applications and (ii) interacts with the conferencing Infrastructure as a Service (IaaS) to allocate virtual conferencing services and associate them to conferencing applications. In order to ensure elastic resource usage and to minimize cost, our proposed PaaS follows an optimal resource provisioning schedule. We formulate this problem and provide an algorithm to solve it.

- ***A Use Case on Network Function Virtualization for the Rapid Provisioning of new Video Delivery Services in Content Delivery Networks***

  Narjes Tahghigh, Sandhya Shanmugasundaram, Jagruti Sahoo, Sami Yangui, **Roch Glitho,** (Université Concordia University)

  A Content Delivery Network (CDN) is an effective solution for offering content services like video on demand to a large number of geographically distributed users. Network Function Virtualization (NFV) enables the rapid provisioning of new services through Virtual Network Functions (VNF) which are loosely coupled software components. This poster provides a use case on video overlay (e.g. adding short video clips as overlays) enhanced by transcoding and compressing. A business model, system architecture and implementation architecture are proposed.

- ***From Relations to Multi-Dimensional Maps: Towards A SQL-to-HBase Transformation Methodology***

  Diego Serrano, Dan Han**, Eleni Stroulia**, (University of Alberta)

  In this paper, we describe a method for transforming and migrating data schemas developed for RDBMS to HBase. The method consists of a set of HBase-organization guidelines and a four-step data-schema transformation process that HBase application developers may follow during the migration of their application data from RDBMS to HBase. The method also considers data-access paths extracted from query logs, in order to improve the quality of the transformation and the eventual access efficiency of the HBase repository. We illustrate and validate the method with a case study.

- ***A Cloud-Based Architecture for Real-Time Multi-modal Interactions and Analytics (Towards the Kaleidoscope App)***

  Hu Zhang, Diego Serrano, **Eleni Stroulia**, (University of Alberta)

  Mobile video streaming becomes increasingly useful in a variety of contexts (social interaction, education and entertainment) and increasingly feasible with the rapid development of wireless networks and mobile technologies. In this project, we develop a platform for multimedia streaming on mobile devices, enhanced with textual and touch-display interactions for a rich user experience. Users (senders) can use our Kaleidoscope mobile application to setup a streaming channel on the platform server, and invite their contacts (receivers) to share their real-time video recordings.

  At the Sender site, the Kaleidoscope app captures the video and shares is with the streaming server. At the Receiver site, the Kaleidoscope app replays the video. At both sites, users can send text messages to the connected peers and touch the display to point out interesting video scenes; the Kaleidoscope app shares these interactions with the whole set of peers. The data (audio/video, text, touch events) will be stored on the cloud server with timestamps to support feature extraction and analytics services on the cloud.

## ➢ POSTER & DEMONSTATION:

- ***Multi-Party Video Conferencing with Bellini on the SAVI Testbed***

  Shuhao Liu, Yinan Liu, **Baochun Li,** (University of Toronto)

SAVI smart applications typically come with requirements on Quality of Experience, which demands speedy data transfers with low latencies. Online video conferencing is one of the representative smart applications with high resource demands. In this demo, we will introduce how Bellini, our efficient, highly flexible, and customizable software-defined application framework can be deployed to support video conferencing applications on the SAVI testbed.

The video conferencing applications are served by multiple Bellini servers disseminated in different SAVI regions. Each user is communicating with an assigned server for the sake of performance. Bellini can automatically deliver the video streams to the desired users via high-performance SAVI inter-regional networks. At the same time, it features flexible control over the traffic, enabling multicast and multi-path forwarding to improve the Quality of Experience.

In this demo, we will show a sample deployment of a high-quality video conferencing application on the SAVI testbed with Bellini.

- ***Live Video Broadcast with Bellini on the SAVI Testbed***

Shuhao Liu, Yinan Liu, **Baochun Li,** (University of Toronto)

Live video broadcast is one of the SAVI smart applications with demands on low-latency bulk data transfers. In this demo, we introduce how Bellini, our efficient, highly flexible, and customizable software-defined application framework can be deployed to support live video broadcast applications on the SAVI testbed.

Bellini powers live video broadcast applications with a number of salient features. First, it minimizes the video playback delays and maximizes the streaming rate by utilizing the high-performance SAVI inter-region network. Second, it supports flexible manipulation over the streaming traffic. For example, application-layer multicast is supported, which minimizes the bandwidth requirements. Third, publishing and subscribing a video stream can be done via standard protocols such as HTTP and RTMP. Various video broadcast clients, such as iOS applications and web browsers, are exhibited in this demo.

- ***Kaleidoscope --- Applications for Multi-Tier Cloud, Smart Edge/Core and Adaptive Resource***

Ronald Desmarais, Andreas Bergen, **Hausi Müller, Sudhakar Ganti,** (University of Victoria)

Imagine the Blue Jays are in 2020 World Series. Reporters comment on the game through live audio and multi-angle video streaming feeds followed by the Rogers Centre audience on their 10G tablets and smartphones. Spectators contribute their own perspectives on the events in the stadium on shared digital billboards. Customized fast food adverts are sent spectators based on their location. Audio and video streams are tagged in real time with timestamps, location, and by subject to provide value-added features such as follow your favourite player, experience that fly ball flying right at you or the view from behind home plate. That is the Kaleidoscope experience.

Using the SAVI infrastructure we are able to dynamically customize, on the fly, user generated video content for further streaming to other users within the same, or different, geographical location. Video is the dominant and fastest growing segment of internet traffic today. Combining user generated video with third party content, contextual information and preference based customizations is possible within SAVI. We utilize EDGE and CORE computing resources of SAVI to achieve a dynamically generated video with customized multiplexed content for each user. We

report on our implementation towards the realization of Kaleidoscope using the SAVI network and integrated smart applications.

## ➢ DEMONSTRATIONS:

### ▪ *The Ignite Distributed Collaborative Scientific Visualization System*

Sushil Bhojwani, Matt Hemmings, Rick McGeer, Yvonne Coady, Ulrike Stege, (University of Victoria), Glenn Ricart, (US Ignite), David Lary, (University of Texas at Dallas), Jens Lincke, (Hasso-Plattner Institute), Dan Ingalls, Robert Krahn, Marko Roeder, (Communication Design Group, SAP)

We describe the Ignite Distributed Collaborative Scientific Visualization System, a system designed to permit real-time interaction and visual collaboration around large data sets, with an initial emphasis on scientific data. The Ignite Visualization System offers such a collaborative environment with real-time interaction on any device between users separated across the wide area.

The Ignite Visualization System provides seamless interaction and immediate updates even under heavy load and when users are widely separated: the design goal was to fetch a data set consisting of 30,000 points from a server and render it within 150 milliseconds, for a user anywhere in the world, and reflect changes made by a user in one location to all other users within a bound provided by network latency. The system was demonstrated successfully on a significant worldwide air pollution data set, with pollution values on a 10km, 25km, 50km, and 100km worldwide grid, with monthly values over an 18-year period. It was demonstrated on a wide variety of clients, including laptop, tablet, and smartphone.

## THEME 2 - RESOURCE CONTROL & MANAGEMENT

## ➢ POSTERS:

### ▪ *EdgeX: Edge Replication for Web Applications*

Hemant Saxena, **Ken Salem,** (University of Waterloo)

Global Web applications face the problem of high network latency due to their need to communicate with distant data centers. Many applications use edge networks for caching images, CSS, javascript, and other static content in order to avoid some of this network latency. However, for updates and for anything other than static content, communication with the data center is still required, and can dominate application request latencies. One way to address this problem is to push more of the web application, as well the database on which it depends, from the remote data center towards the edge of the network. In this poster, we present preliminary work in this direction. Specifically, we present an edge-aware dynamic data replication architecture for relational database systems supporting web applications. Our objective is to allow dynamic content to be served from the edge of the network, with low latency.

### ▪ *Self Partitioning, Edge-Aware Replicated DataBase*

Cătălin Avram and **Ken Salem**, (University of Waterloo)

SpearDB is a Self-Partitioning, Edge-Aware Replicated DataBase. It provides OLTP capabilities; ACID guarantees and executes transactions under Parallel Snapshot Isolation (PSI). The main contributions of this work are a dynamic, automatic data-partitioning algorithm and an edge-aware implementation of PSI. The former allows for transparent usage of a distributed database without the need to specify partitions *a priori*, while at the same time optimizing data placement to improve transactional throughput for a dynamically changing workload. The latter takes advantage of the central position of the core to lower network utilization and introduce new ways of committing transactions, new live data migration capabilities and stored procedures - all under PSI guarantees.

- ***Lightweight Robust Optimizer for Distributed Application Deployment in Multi-Clouds***

  Ravneet Kaur, **Murray Woodside, John Chinneck** (Carleton University)

  A critical challenge to distributed applications in multi-clouds is response latency due to communications delays between the components.

  Acceptable latency requires careful deployment, and rapid adaptation requires a lightweight analysis. A lightweight optimizer has been designed to partition components between an edge and a core cloud which satisfies a hard constraint on response latency, capacity constraints in both clouds, and which minimizes power consumption to the extent possible. The optimizer partitions the communications graph for the components using graph coarsening and refining, as in the Kernighan-Lin-Verbelen approach, combined with bin packing. Experiments show it succeeds much better than competing approaches, judged on execution time and the ability to meet the response constraint.

- ***Evaluating Cluster Configurations for Big Data Processing: An Exploratory Study***

  Roni Sandel, Marios Fokaefs, Mark Shtern, **Marin Litoiu,** (York University)

  In this work, we explore the challenges around evaluating the performance of various Big Data cluster configurations in order to decide the optimal one. We perform a series of experiments with various configurations focusing on three factors that may affect the cluster's response time: data compression, data schema and cluster topology. Eventually, the outcomes of our study are encapsulated in a performance model that predicts the cluster's response time as a function of the incoming workload and evaluates the cluster's performance less costly and faster. This systematic and effortless evaluation method will facilitate the selection and migration to a better cluster as the performance and budget goals change. We use HBase as the large data processing cluster and we conduct our experiments on traffic data from a large city and on a distributed community cloud infrastructure.

- ***Engineering Economics of Everything-as-a-Service (XaaS)***

  Marios Fokaefs and **Marin Litoiu**, (York University)

  Web and mobile technologies along with the virtualization of computation and storage resources have allowed extensive flexibility in the development, deployment and delivery of software solutions. This flexibility has given rise to new business models and interesting implications in the economics of software ecosystems. As in every economic model, the two fundamental dimensions are the value that is created by software and flows in the ecosystem and the cost to reserve resources (human, computational and other). In this work, we discuss the principles around the creation of economic models to calculate value and cost of software. The goals of such models

would be to estimate profits, evaluate alternative options and also predict the economic state of software ecosystems under certain conditions. Finally, we demonstrate the use of such models in example ecosystems to assist the decision-making process of software stakeholders.

## ➢ POSTER & DEMONSTRATION:

### ▪ *OpenFlow-enabled DDoS Attack Mitigation on Hybrid Cloud*

Nasim Beigi-Mohammadi, Cornel Barna, Hamzeh Khazei, **Marin Litoiu**, (York University)

Application layer Distributed Denial of Service (DDoS) attacks are becoming serious concerns as multiple attack machines are harder to detect than one attack machine. Moreover, the behavior of each attack machine can be stealthier, making it more difficult to track and shut down. Therefore, mitigating such attacks is of essential importance. In this work, we design and implement a DDoS attack mitigation solution that protects application clusters on a public cloud by leveraging SAVI cloud and SDNLauncher. The mitigation solution isolates the suspicious traffic from the regular application clusters through redirecting the suspicious traffic to SAVI. This enables us to exercise control and perform closer surveillance of the suspicious traffic. Thus, on one hand, the application clusters are protected against the attack, and on the other hand, further analyses are performed over the isolated traffic to differentiate between malicious and non-malicious traffic. We show how OpenFlow and SDN help to take advantage of public and private clouds in building a hybrid and adaptive mitigation solution against DDoS attacks.

### ▪ *K-Feed, A Data-Oriented Approach to Application Performance Management in Cloud*

Saeed Zareian, Rodrigo Veleda, Hamzeh Khazei, **Marin Litoiu,** (York University)

This poster presents K-Feed (Knowledge Feed), a platform for real-time application performance analysis and provisioning in the cloud. K-Feed can perform at-scale monitoring, analysis, and provisioning of cloud applications. We explain the components, the implementation and the validation of the K-feed platform. To illustrate its feasibility, we use it to monitor and build a performance model of a clustered web application. To model the application, we use off-the-shelf components.

## ➢ DEMONSTRATIONS:

### ▪ *Apache Spark on SAVI with OpenStack Sahara*

Hamzeh Khazaei, Marios Fokaefs, **Marin Litoiu,** (York University)

This demo aims to demonstrate the synergy between the OpenStack Sahara service and Apache Spark OpenStack Sahara service and Apache Spark to set up data-intensive application clusters in the SAVI cloud. We will first show you how to set up a Sahara/Spark cluster by creating the individual components including the cluster template and the node group templates (for the master and slave node groups) and, finally, the cluster itself. Then, we will show you how to develop and deploy Spark jobs on the deployed Sahara cluster and see the processing results. The Spark jobs will analyze traffic data from the Connected Vehicles and Smart Transportation (CVST) project. Finally, we will show you how to scale up the cluster, in case of high workload.

- ***Enabling an Enhanced Data-as-a-Service Ecosystem***

  Justin Cuaresma, Ethan Nguyen, Mark Shtern, **Marin Litoiu,** (York University)

  Data-as-a-service is a concept that enables providers to expose proprietary data; it also controls how interested users consume data. With the advent of cloud and virtual infrastructure, the concept of enhanced data-as-a-service has emerged, where the provider actually offers a combination of platform-as-a-service and data-as-a-service. In this combination, the users are able to perform analytic tasks, local to where the data resides, whilst reducing the (expensive and slow) transmission of data over networks. This prototype demonstrates an ecosystem that allows providers to precisely share portions of their data with users, using a model where users submit analytic jobs that run on the provider's big data infrastructure.

- ***Hogna: A Platform for Developing Self-Adaptive Applications in Cloud***

  Carnel Barna and **Marin Litoiu,** (York University)

  Data-as-a-service is a concept that enables providers to expose proprietary data; it also controls how interested users consume data. With the advent of cloud and virtual infrastructure, the concept of enhanced data-as-a-service has emerged, where the provider actually offers a combination of platform-as-a-service and data-as-a-service. In this combination, the users are able analytic tasks, local to where the data resides, and reducing the (expensive and slow) transmission of data over networks. This prototype demonstrates an ecosystem that allows providers to precisely share portions of their data with users, using a model where users submit analytic jobs that run on the provider's big data infrastructure.

## THEME 3 - SMART CONVERGED EDGE

➢ **POSTERS:**

- ***Robust Resource Reservation in Virtual Wireless Networks***

  Ali Abbasi and **Majid Ghaderi**, (University of Calgary)

  In this work, we study resource reservation in virtual wireless networks with the aim of minimizing the operational cost of the network. With this regard, the main constraint facing the operator is that only limited information about future traffic demand is typically available to the operator. To address this issue, we investigate reservation policies that are robust to the worst-case traffic demand that fit the available information i.e., the policies that minimize the worst-case expected operational cost. The problem is formulated for several resource reservation options that are offered commonly in practice. For each case, convexity of the problem is discussed and then its dual form is presented as a semidefinite program. While, semi definite programs can be solved in polynomial time, the optimal closed form reservation policies are obtained for several practical cases. Moreover, the worst-case cost of these policies are analytically compared to the expected cost of the algorithm that has full knowledge of the future demand.

- ***Simple Distributed Programming in Software-Defined Networks***

  Soheil Hassas Yeganeh, **Yashar Ganjali,** (University of Toronto)

Simplicity of network programming is a voiced advantage of Software-Defined Networking (SDN). In networks of larger scale, however, developing a distributed control application demands significant efforts to design, optimize, and maintain in the absence of proper platform-level support.

We present the design and implementation of Beehive, a distributed control platform that simplifies distributed programming in SDN. In Beehive, control applications are centralized asynchronous message handlers that optionally store their state in dictionaries (i.e., generic associative arrays or maps). This enables Beehive's control platform to automatically infer how an application accesses its state in response to messages, and allows a seamless transformation of the application into a distributed system while preserving the application's intended behavior. In contrast to existing distributed controllers, Beehive hides the common boilerplates of distributed programming, such as coordination, locking, concurrency, consistency and fault-tolerance. More importantly, Beehive is capable of runtime instrumentation of control applications. This enables Beehive to live migrate the workload among controllers aiming to optimize the control plane as a whole. The runtime instrumentation, as a feedback provided to application developers, can be used to identify bottlenecks in control application design, and can help developers enhance their designs. We evaluate our system by implementing several use-cases.

- ***Dynamic Proactive Routing in Software Defined Networks***

Sajad Shirali-Shahreza and **Yashar Ganjali** (University of Toronto)

The ability to manage individual flows is a major benefit of Software-Defined Networking. The overheads of this fine-grained control, e.g. initial flow setup delay, can overcome the benefits, for example when we have many time-sensitive short flows. The flow setup delay is one of the main side effects of individually managing flows, which is also known at reactive routing. In this work, we designed a dynamic proactive routing that eliminates flow setup delay. Our dynamic proactive routing uses wildcard rules to manage groups of flows, and our recently proposed FleXam sampling extension to OpenFlow to collect information about individual flows matching each installed wildcard rules. It will monitor traffic of all installed rules and whenever one of them becomes too large, it will break it down into multiple rules and route them individually to distribute rules and prevent congestion. Our simulation result shows that this will significantly reduce flow termination time for small flows, while reduces the table occupancy and controller load.

- ***Dynamic Resource Allocation for MapReduce with Partitioning Skew***

Zhihong Liu (Univ. of Waterloo), Qi Zhang (Univ. of Waterloo & Toronto), **Raouf Boutaba** (Univ. of Waterloo), Yaping Liu and Zhenghu Gong (National University of Defense Technology)

MapReduce has become a popular model for large-scale data processing in recent years. However, existing MapReduce schedulers still suffer from an issue known as partitioning skew, where the output of map tasks is unevenly distributed among the reduce tasks. In this poster, we present DREAMS, a framework that provides runtime partitioning skew mitigation. Unlike previous approaches that try to balance the workload of reducers by repartitioning the intermediate data assigned to each reduce task, in DREAMS we cope with partitioning skew by adjusting task runtime resource allocation. We show that our approach allows DREAMS to eliminate the overhead of data repartitioning. Through experiments using both real and synthetic workloads running on a 11node virtual Hadoop cluster, we show that DREAMS can effectively mitigate negative impact of partitioning skew, thereby improving job performance by up to 20:3%.

- ***SiMPLE: Survivability in MultiPath***

  Md Mashrur Alam Khan, Nashid Shahriar, Reaz Ahmed, **Raouf Boutaba,** (University of Waterloo)

  Internet applications are deployed on the same network infrastructure, yet they have diverse performance and functional requirements. The Internet was not originally designed to support the diversity of current applications. Network Virtualization enables heterogeneous applications and network architectures to coexist without interference on the same infrastructure. Embedding a Virtual Network (VN) into a physical network is a fundamental problem in Network Virtualization. A VN Embedding that aims to survive physical (e.g., link) failures is known as the Survivable Virtual Network Embedding (SVNE). One way of ensuring survivability is to allocate redundant resources during the embedding process.

  A key challenge in the SVNE problem is to ensure VN survivability with minimal resource redundancy. To address this challenge, we propose SiMPLE. By exploiting path diversity in the physical network, SiMPLE provides maximal VN survivability while incurring minimal resource redundancy. We formulate this problem as an ILP and implement it using GLPK. We also propose a greedy algorithm to solve the larger instances of the problem. Simulation results show that SiMPLE outperforms full backup and shared backup schemes for SVNE, and produces nearoptimal results.

- ***nf.io: A File System Abstraction for NFV Orchestration***

  Md. Faizul Bari, Shihabur Rahman Chowdhury, Reaz Ahmed, **Raouf Boutaba**, (University of Waterloo)

  In recent years, Network Function Virtualization (NFV) has gained a lot of traction from both industry and academia. NFV promotes vendor independence and rapid evolution through open source software, open standards and open APIs. However, adopting these principles for virtual middleboxes or Virtual Network Functions (VNFs) is not enough. The VNF management and orchestration systems also need to adopt the same principles, otherwise a network operator may still face vendor lockin. Moreover, standardization efforts take a long time to converge and are often futile. For this reason, we introduce nf.io that uses existing well known Linux file system interface for VNF management and orchestration. We have developed a prototype, and provided some example usecases to demonstrate its effectiveness.

- ***Aurora: Adaptive Block Replication in Distributed File Systems***

  Qi Zhang, Sai Qian Zhang, **Alberto Leon-Garcia**, University of Toronto, **Raouf Boutaba**, (University of Waterloo)

  Distributed file systems such as the Google File System and Hadoop Distributed File System have been used to store large volumes of data in Cloud data centers. These systems divide data sets in blocks of fixed size and replicate them over multiple machines to achieve both reliability and efficiency. Recent studies have shown that data blocks tend to have a wide disparity in data popularity. In this context, the naïve block replication schemes used by these systems often cause an uneven load distribution across machines, which reduces the overall I/O throughput of the system. While many replication algorithms have been proposed, existing solutions have not carefully studied the placement of data blocks that balances the load across machines, while ensuring node and rack level reliability requirements are satisfied. In this poster, we study the dynamic data replication problem with the goal of balancing machine load while ensuring machine and rack level reliability requirements are met.

We propose several local search algorithms that provide constant approximation guarantees, yet simple and practical for implementation. We further present Aurora, a dynamic block placement mechanism that implements these algorithms in the Hadoop Distributed File System with minimal overhead. Through experiments using workload traces from Yahoo! and Facebook, we show Aurora reduces machine load imbalance by up to 26:9% compared to existing solutions, while satisfying node and racklevel reliability requirements.

- ***Network Function Virtualization Enabled Multicast Routing on SDN***

Sai Qian Zhang, Qi Zhang, Hadi Bannazadeh, **Alberto Leon-Garcia,** (University of Toronto)

Large-scale online services today often span multiple geographically distributed domains and require significant storage, compute and networking resources. Furthermore, these services often need to be scaled up and down dynamically according to service demand fluctuation. To facilitate efficient resource allocation while allowing multiple online services to share to physical hosting infrastructure (e.g. data centers), there is a emerging trend towards allocating resources to service applications in the form of virtual infrastructures (VI) that consist of virtual machines, virtual routers and switches interconnected by virtual links. However, despite extensive study of virtual infrastructure scheduling algorithms, existing work has not studied the problem of provisioning VIs dynamically according to demand fluctuations. To address limitation, in this work we present a framework for virtual infrastructure provisioning that adjusts virtual infrastructure resource allocation according to demand fluctuations, while satisfying performance requirements. We demonstrate the effectiveness of our framework through simulations in realistic application scenarios.

- ***Fast Network Flow Resumption for Live Virtual Machine Migration on SDN***

Sai Qian Zhang, Pouya Yasrebi, Ali Tizghadam, Hadi Bannazadeh, **Alberto Leon-Garcia,** (University of Toronto)

Virtual machine (VM) migration occurs very frequently in the communication networks. VM Migration enables a running OS, including memory and storage to move from one physical host to another physical host. A particular case of interest is live migration where the process of migrating the full state from one OS to the other should happen continuously and without any connection disruption. In order to have a seamless VM migration process the system has to be able to resume network connectivity in a very fast pace. Fast resumption is proved to be a challenging problem. In this paper, we present a routing scheme to efficiently migrate the flow of the VM by reusing the existing rules in the OpenFlow enabled switches to achieve fast network flow resumption. We formulate the problem by a integer programming problem, we prove its NP-completeness and a heuristic algorithm is proposed to solve the problem. Software simulation and real testbed implementation are done to demonstrate the performance of the migration scheme.

- ***Anomaly Detection and Localization in Software Defined Infrastructures***

Qi Zhang, Jieyu (Eric) Lin, Hadi Bannazadeh, **Alberto Leon-Garcia,** (University of Toronto)

Software-Defined Infrastructure (SDI) provides a unified framework for managing heterogeneous virtualized resources in Cloud infrastructures. A key challenge in the realization of SDI is detection and localization of abnormal network and system behaviors. In this work, we present our technical solution to the anomaly detection problem in SDI. We first identify abnormal behavior of individual components based on their characteristics, and then provide a technique that effectively localizes

the anomalies for better analysis and diagnosis. We demonstrate the benefit of our approach through several realistic test scenarios on the SAVI testbed.

## ➢ POSTER & DEMONSTRATION:

### ▪ *Multicast as a Service on SAVI Network*

Sai Qian Zhang, Hadi Bannazadeh, **Alberto Leon-Garcia**, (University of Toronto)

Group communication is widely used in computer network systems. When more than one receiver is involved in the data transmission mechanism, multi-cast is the most efficient and reliable solution. In this project, we propose an efficient multicast algorithm with the goal of saving bandwidth and minimizing transmission delay on the SAVI testbed. We have already implemented this multicast mechanism on the SAVI edge network using the API provided by the SDI manager. We also developed the user interface to facilitate the multicast service.

# THEME 4 - INTEGRATED WIRELESS/OPTICAL ACCESS

## ➢ POSTERS:

### ▪ *Leveraging Synergy of SDWN and Multi-Layer Resource Management for Network Optimization*

Mahsa Derakhshani, Saeedeh Parsaeefard, **Tho Le-Ngoc**, (McGill University), **Alberto Leon-Garcia**, (University of Toronto)

Convergence of software-defined networking (SDN), software-defined radio (SDR), and virtualization on the concept of software-defined wireless networking (SDWN) supports flexibility to offer diverse services for the next generation of wireless networks. The principal technique behind the service-oriented SDWN is the separation of the control and data planes, from the deep core entities to edge wireless access points. This separation allows the abstraction of resources as transmission parameters of each user over the SDWN. Similar to traditional wireless networks, in this user-centric, service-oriented and integrated environment, resource management plays a critical role to achieve efficiency and reliability. In this work, a simple, efficient and integrated structure for SDWN is studied based on its three pillars: SDN, SDR, and virtualization. For this networking paradigm, we provide a general architecture and discuss that the SDWN structure enables converged multi-layer resource management, while we believe that CML resource management enabled by SDWN can considerably increase the network performance.

### ▪ *Virtualization of Multi-Cell 802.11 Networks: Association and Airtime Control*

Mahsa Derakhshani, Xiaowei Wang, **Tho Le-Ngoc,** (McGill University), **Alberto Leon-Garcia,** (University of Toronto)

This work investigates the station (STA)-access point (AP) association and airtime control for virtualized 802.11 networks to provide service customization and fairness across multiple Internet Service Providers (ISPs) sharing the common physical infrastructure and network capacity. More specifically, an optimization problem is formulated on the STAs' transmission probabilities to maximize the overall network throughput, while providing airtime usage guarantees for the ISPs. An

algorithm to reach the optimal solution is developed by applying monomial approximation and geometric programming iteratively. Based on the proposed three-dimensional Markov chain model of the enhanced distributed channel access (EDCA) protocol, the detailed implementation of the optimal transmission probability is also discussed. The accuracy of the proposed Markov-chain model and the performance of the developed association and airtime control scheme are evaluated through numerical results.

- ***Adaptive Pilot Duration Allocation in Wireless Virtualized Networks with Massive MIMO***

Rajesh Dawadi, Saeedeh Parsaeefard, Mahsa Derakhshani, **Tho Le-Ngoc**, (McGill University)

In this work, a resource allocation problem for a wireless virtualized network (WVN) is investigated where each base station (BS) is equipped with a large number of antennas, called massive MIMO. We focus on the scenario in which, due to the pilot contamination error, the perfect estimation of channel state information (CSI) of users is not available. In this case, the duration of pilot sequence transmission plays a critical role on the achieved total throughput of WVN. Therefore, we consider this parameter as a new optimization variable and propose a novel utility function for the resource allocation formulation. In order to solve the resource allocation problem, we propose a two-step iterative algorithm: 1) at the first step, one set of resource allocation parameters, including power, sub-carrier and number of antennas, are derived; 2) at the second step, the pilot duration is optimized. Simulation results reveal that properly designing the pilot duration improves the WVN performance compared to the case of fixed pilot duration.

- ***Full Duplex Analog WiFi over PONs***

Zhihui Cao, An T. Nguyen, **Leslie A. Rusch,** (Université Laval)

We experimentally demonstrated the transparent convergence of distributed antenna system in a digital wavelength division multiplexed passive optical network system employing reflective semiconductor optical amplifier (RSOA) at the remote antenna unit (RAU). By incorporating the up/down mixer in Central Office and RAU, we experimentally achieve full-duplex WiFi transmission along with 1 Gb/s downlink digital on-off-keying signal in single mode fiber up to 15 km (64QAM), 40 km (16QAM), and 60 km (QPSK) with a 1 meter wireless path.

- ***Virtualized WiFi on Radio Over Fiber for PONs***

Zhihui Cao, Robert Morawski, Quang Dong Ho, Thanh Ngon Tran, **Tho Le-Ngoc**, (McGill University), **Leslie Rusch**, (Unversité Laval)

The radio over fiber experimental test bed at UL will be combined with the software radio platform from McGill Univ. to complete a proof of concept demonstration. Full duplex operation at the optical testbed allows real-time testing of virtualization of wireless resources at the smart edge of the network.

## ➢ POSTER & DEMONSTRATION:

- ***Virtualized and Software-Defined WiFi Radio Access Networks***

Quang-Dung Ho, Alfred Kenny, **Tho Le-Ngoc**, (McGill University)

Wireless virtualization and software-defined radio access networking (S-RAN) are key technologies in next-generation communications networks. They make the networks more efficient in terms of resource utilization and power consumption, less expensive, and more flexible in technological evolution. This work designs and implements a centrally managed, software-defined and virtualized IEEE 802.11/WiFi access network. Flow-based and protocol-based virtualization techniques are focused. Multiple virtual access points (vAPs) can be dynamically created or deleted from a single physical AP. Each vAP has its own basic service set (BSSID), data path and wireless resource. Software-defined IEEE 802.11/WiFi radios are developed to support protocol-based virtualization.

## THEME 5 - APPLICATION PLATFORM TESTBED

### ➢ POSTER:

- ***Expanding OpenFlow Capabilities with Virtualized Reconfigurable Hardware***

  Stuart Byma, Naif Tarafdar, Talia Xu, Hadi Bannazadeh, **Alberto Leon-Garcia, Paul Chow,** (University of Toronto)

  We present a novel method of using cloud-based virtualized reconfigurable hardware to enhance the functionality of OpenFlow Software-Defined Networks. OpenFlow is a capable and popular SDN implementation, but when users require new or unsupported packet-processing, software processing in the OpenFlow controller cannot provide multi-gigabit rates. Our method sees packet flows redirected through engines that can add new capabilities to an OpenFLow network, while retaining line-rate processing. A case study shows this can be achieved with virtually no loss in throughput and minimal latency overheads.

- ***OpenCLoud: A Heterogenous Reconfigurable Cloud Platform with an OpenCL Front-end***

  Naif Tarafdar, Eric Fukuda, William Suriaputra, **Paul Chow,** (University of Toronto)

  We present OpenCLoud; this is a virtualized heterogenous cloud platform. This heterogenous platform includes FPGAs, GPUs and CPUs. The users, using an OpenCL application can interface with cloud management software to build their own platform of virtualized devices which they can use for their applications. Furthermore multiple users can share these devices from the same cloud network without being aware of this. These intricacies are hidden from the user through the use of standard OpenCL API. Several case-studies have been investigated including image processing application, financial applications and query applications.

- ***Interactive Large-scale Data Analysis on Virtual FPGAs in the Cloud***

  Eric Fukuda, Naif Tarafdar, **Paul Chow,** (University of Toronto)

  As more and more data is gathered and stored in the cloud, the computation required to analyze it is increasing. BigQuery is a Google service that enables users to analyze large-scale data stored in the cloud with SQL queries on the fly by distributing the computation over thousands of servers with software processors. We propose a system for improving the power efficiency of such a system by using multiple FPGAs distributed in the cloud. FPGAs are known to improve the efficiency of SQL query processing, and database appliances that use FPGAs have been developed by several computer makers. However, accelerating interactive query processing such as BigQuery on a

heterogeneous cloud platform presents further challenges, e.g., real-time allocation and sharing of FPGAs and network routing among distributed computation nodes. We employ an FPGA virtualization technique to overcome these challenges and evaluate the effectiveness of our method with Apache Drill, an open source implementation of BigQuery.

- ***SDI Manager Architecture***

Byungchul Park, Thomas Lin, Hadi Bannazadeh, **Alberto Leon-Garcia,** (University of Toronto)

Software Defined Infrastructure (SDI) enables programmability of infrastructure by enabling the support of cloud-based applications, customized network functions, and hybrid combinations of these. The SDI manager provides integrated resource management for converged heterogeneous resources by abstraction in SDI environment. In this work, we present the SDI manager architecture that supports multiple OpenFlow controllers, QoS, and multiple OpenFlow protocol versions.

- ***Enhancing VoIP Data Plane using SDN***

Lilin Zhang, Ali Tizghadam, Hadi Bannazadeh, **Alberto Leon-Garcia,** (University of Toronto)

The quality of VoIP service heavily replies on the end-to-end network condition. Media sessions that are routed through highly-utilized links are in more jeopardy of longer delays, bigger packet loss rate, and thus receive poor user perceived quality. Changing the packet routes of degraded sessions can improve the media quality, but is not an easy operation to do in IP networks. SDN provides the unprecedented centralized solution in such an occasion: the embedding of flows in the data plane, once calculated, can be quickly deployed via the SDN controller.

- ***Operations, Maintenance and Diagnostics for SAVI Testbed SDI***

Joseph Wahba, Hadi Bannazadeh, **Alberto Leon-Garcia**, (University of Toronto)

This poster describes the different set of tools that we use in our SAVI testbed in order to perform Operations, Maintenance and Diagnostics. We used three main categories of tools which are Nagios servers, Automated Bash Testing Scripts and logical tests such as Janus Network Flow diagnostics and Whale Link Diagnostics. The previously mentioned tools help SAVI testbed administrators monitor the system and detect any faults and act accordingly.

- ***Idle VM Consolidation in SAVI Testbed***

SeyedAli Jokar, Joseph Wahba, Hadi Bannazadeh, **Alberto Leon-Garcia**, (University of Toronto)

With the increasing demand and interest in cloud service, new methods were introduced to define the architecture of the infrastructure as a service. The SAVI test-bed controls and manages converged virtual resources enforced with the power of the software-defined networking paradigm, namely the Software-Defined Infrastructure (SDI). SAVI uses OpenStack to perform Orchestration. OpenStack sets a limit on the number of virtual machines since the test-bed has limited number of hardware resources. At any point of time there are some idle VMs which under-utilize their dedicated hardware resources. Hence, dedicated hardware resources of those idle VMs could be multiplexed by using higher sharing factors. Moreover, we will show how consolidating those idle VMs would affect the performance. SAVI Resource Allocation (SAVIRA) algorithm diagnoses and migrates idle VMs based on the owner's Service Level Agreement (SLA) of the VM and its CPU utilization. In this poster we have theoretical proof that migration of up to 83 VMs will degrade the performance by 1% in the worst case scenario.

- ***Security Function Virtualization in Software Defined Infrastructure***

  Pouya Yasrebi, Sina Monfared, Hadi Bannazadeh, **Alberto Leon-Garcia**, (University of Toronto)

  In this work we have designed an approach to implement security as a Virtualized Network Function (VNF) that can be implemented on a Software-Defined Infrastructure (SDI). We present a scalable, flexible, and seamless design for a deep packet inspection (DPI) system for network intrusion detection and prevention. We discuss how our design introduces significant reductions in both capital and operational expenses (CAPEX and OPEX). As proof of concept, we describe an implementation for a modular security solution that uses the SAVI SDI testbed. We discuss our testing methodology and provide measurement results for the test cases where a cloud faces various security attacks.

- ***Scalability in Security Function Virtualization of Software Defined Infrastructure***

  Pouya Yasrebi, Sai Qian Zhang, Hadi Bannazadeh, **Alberto Leon-Garcia**, (University of Toronto)

  In this work we offer multiple dimensions for scalability of security as a Virtualized Network Function (VNF). In our previous work we designed an approach to implement security as a VNF. Both Intrusion Detection System (IDS) and Intrusion Prevention System (IPS) in our design have been studied and we removed some of the bottlenecks in this work. We also discuss a damage reduction problem with network constraints for IPS and provide solutions in different scenarios. As proof of concept, we implement an attack scenario and use our algorithm in a simple case to reduce damage from the attackers to the webservers that use the SAVI SDI testbed. We discuss our testing methodology and provide measurement results for the test case.

- ***SAVI-GENI Cloud Federation: Resource Procurement***

  Sushil Bhojwani, Riz Panjwani, Rick McGeer, Sudhakar Ganti, Andi Bergen, Ron Desmarais, **Hausi Müller** (University of Victoria), Hadi Bannazadeh (University of Toronto)

  GENI and SAVI systems are both modified versions of the OpenStack Cloud system. In this work, both the GENI User and the SAVI User have been set up with the necessary tools and credentials in order to access GENI or SAVI resources using the Common Interface. Once the users have been setup with the Common Interface, they can request resources on either GENI or SAVI cloud by issuing various commands. The users can also query the resources available or allocated among other things. The commands issued by the user are translated by the Common Interface to those that can be understood by a corresponding client for either GENI or SAVI cloud. For the GENI cloud, it is the OMNI Command Client. For the SAVI cloud, it is the NOVA Command Client. The NOVA or OMNI clients then execute these commands against their appropriate clouds. The responses are parsed and displayed to the user via the Common Interface. The main module that we worked on was the Common Interface module. We have covered the authentication process in a separate poster.

## ➢ POSTER & DEMONSTRATION:

- ***MonArch Multilayer Monitoring and Analytics in Software-Defined Infrastructures***

  Jieyu (Eric) Lin, Hadi Bannazadeh, **Alberto Leon-Garcia,** (University of Toronto)

Software-Defined Infrastructure provides a unified architecture for integrated management of virtualized heterogeneous resources in cloud infrastructures. Monitoring and measurement is crucial for enabling intelligent management in the Software-Defined Infrastructure. Flexible scalable analytics and Monitoring-as-a-Service (MaaS) are two key features of next generation monitoring systems. To meet these requirements, we designed and implemented MonArch, a cross-layer MaaS system that provides collection, storage, and analytics capabilities. MonArch has been deployed and operational in the SAVI Testbed.

- ***Software Defined Infrastructure overlay over heterogeneous cloud to facilitate NFV service chaining***

  Spandan Bemby, SeyedAli Jokar, Hadi Bannazadeh, **Alberto Leon-Garcia**, (University of Toronto)

  Software defined infrastructure (SDI) is a means for integrated control and management of converged heterogeneous resources. SDI offers many benefits such as greater flexibility by facilitating dynamic NFV service chaining. However these benefits are restricted to infrastructures that have SDI capabilities. We propose to bring these capabilities to non SDI infrastructures, or infrastructures where these capabilities may not be fully exposed to end users. We achieve this by leveraging overlay networks and managing them in the user space. Specifically, we demonstrate this "overlay SDI management" in the context of dynamic NFV service chaining.

- ***Flexible NFV orchestration using Containerized Overlay***

  SeyedAli Jokar, Spandan Bemby, Hadi Bannazadeh, **Alberto Leon-Garcia**, (University of Toronto)

  Cloud platforms consist of pools of heterogeneous resources providing virtualized computing resources with limited flexibility and configuration. To address these issues, Software Defined Infrastructure (SDI) enhances traditional networking with Software Defined Networking (SDN) to make a software defined overlay on top of cloud resources. This virtual overlay can be configured and managed like a real network infrastructure. Moreover, integration with container brings multiple benefits including portability, reusability, and rapid deployment to this overlay network in terms of application component. We have used docker to containerize the overlay, to isolate the environment of each application component. In addition, the SDI can set up the communication topology between these containers, if required. The overlay is created by creating tunnels between the VMs and each VM hosts multiple containers. In this poster a network function virtualization scenarios is being discussed as potential use case of this containerized overlay network.

- ***NFV Realization on Virtual Customer Premise Edge (vCPE) using SAVI SDI***

  Sina Monfared, Hadi Bannazadeh, **Alberto Leon-Garcia**, (University of Toronto)

  In this work, we investigate integration of Virtual Customer Premises' Edge (vCPE) using SAVI Software Defined Infrastructure (SDI). In particular, we implemented a proxy application within the home router as a use case of vCPE and Network Function Virtualization (NFV). The proxy is shown to be useful for applications such as website caching and unwanted traffic blocking, including website blacklisting.

- ***SAVI-GENI Cloud Federation: Resource Procurement***

  Sushil Bhojwani, Riz Panjwani, Rick McGeer, **Sudhakar Ganti,** Andi Bergen, Ron Desmarais, **Hausi Müller** (University of Victoria), Hadi Bannazadeh (University of Toronto)

GENI and SAVI systems are both modified versions of the OpenStack Cloud system. In this work, both GENI User and SAVI User have been set up with necessary tools and credentials in order to access GENI or SAVI resources using the Common Interface. Once the users have been setup with the Common Interface, they can request resources on either GENI or SAVI cloud by issuing various commands. The users can also query the resources available or allocated among other things. The commands issued by the user are translated by the Common Interface to those that can be understood by a corresponding client for either GENI or SAVI cloud. For the GENI cloud, it is the OMNI Command Client. For the SAVI cloud, it is the NOVA Command Client. The NOVA or OMNI clients then execute these commands against their appropriate clouds. The responses are parsed and displayed to the user via the Common Interface. The main module that we worked on was the Common Interface module. We have covered the authentication process in a separate poster.

- ***SAVI-GENI Cloud Federation: Authentication***

Sushil Bhojwani, Riz Panjwani, Rick McGeer, **Sudhakar Ganti,** Andi Bergen, Ron Desmarais, **Hausi Müller** (University of Victoria)

GENI and SAVI systems are both modified versions of the OpenStack Cloud system. SAVI and GENI users have different processes for getting authenticated. For a SAVI user to access GENI resources, they must visit the GENI portal from where they get redirected to SAVI Shibboleth Identity Provider (SSIDP). SSIDP verifies the SAVI user against a separate LDAP database of all SAVI users and redirects the user back to GENI portal along with appropriate credentials which allows the GENI Portal to verify the user. Once the user has access to the GENI Portal, they can request resources using the web interface, or they can download an OMNI Configuration Bundle which allows them to setup a local environment in order to secure GENI resources.

We provide a Common Interface tool which they can use to issue commands to GENI in order to secure resources. We cover this process/architecture in a separate document. For a GENI user to access SAVI resources, they must login to the GENI portal which redirects them to SAVI portal along with a certificate which verifies them to SAVI network. SAVI then allows the GENI user to create a SAVI account, which allows them to access SAVI resources using the SAVI web portal, or to access them from a local environment using our Common Interface tool to issue commands to SAVI in order to secure resources as stated above. The main modules that we worked on were the SAVI Shibboleth Identity Provider and the accompanying LDAP database and scripts, as well as the ability for GENI users to secure access to SAVI portal using an X.509 certificate.

## ➢ DEMONSTRATIONS:

- ***SAVI-GENI Cloud Federation***

Sushil Bhojwani, Riz Panjwani, Rick McGeer, **Sudhakar Ganti,** Andi Bergen, Ron Desmarais, **Hausi Müller,** (University of Victoria), Hadi Bannazadeh (University of Toronto)

In this demo, we show the inter working of SAVI-GENI cloud federation. This work involves the tools and techniques developed in the two works related to SAVI-GENI Federation of Resource Procurement and Authentication. This demo shows how a SAVI user requests resources on a GENI cloud and vice versa.

## ACKNOWLEDGEMENT

*This research has been funded in part or completely by the Smart Applications on Virtual Infrastructure (SAVI) project funded under the National Sciences and Engineering Research Council of Canada (NSERC) Strategic Networks grant number NETGP397724-10.*