

“SMART APPLICATIONS ON VIRTUAL INFRASTRUCTURE”

2016 Annual General Meeting Agenda: Wednesday, July 6, 2016 - 8:25 am to 3:15 pm

WELCOME & OVERVIEW	8:25 am	Al Leon-Garcia - SAVI Scientific Director, University of Toronto	
PROGRESS REPORT	8:30 am	Baochun Li - Theme 1 Lead, University of Toronto	
	8:45 am	Hausi Muller - Theme 1, University of Victoria	
	Theme 1: Smart Applications	8:55 am	Eleni Stroulia - Theme 1, University of Alberta
		9:05 am	Roch Glitho - Theme 1, Université Concordia University
	Theme 2: Resource Control & Management	9:15 am	<i>Questions and Answers Period</i>
		9:25 am	Marin Litoiu - Theme 2 Lead, York University
		9:40 am	Murray Woodside - Theme 2, Carleton University
		9:50 am	Ken Salem - Theme 2, University of Waterloo
		10:00 am	<i>Questions and Answers Period</i>
	10:10 am	<i>Coffee Break and Presentation Session 1: Theme 1</i>	
PROGRESS REPORT	11:00 pm	Raouf Boutaba - Theme 3 Lead, University of Waterloo	
	11:15 pm	Yashar Ganjali - Theme 3, University of Toronto	
Theme 3: Smart Converged Edge	11:25 pm	Al Leon-Garcia - Theme 3, University of Toronto	
	11:35 pm	Carey Williamson - Theme 3, University of Calgary	
	11:45 pm	<i>Questions and Answers Period</i>	
Theme 4: Integrated Wireless/Optical Access	11:55 pm	Leslie Rusch - Theme 4 Lead, Université Laval	
	12:10 pm	Tho Le Ngoc - Theme 4, McGill University	
	12:20 pm	<i>Questions and Answers Period</i>	
	12:30 pm	<i>Lunch and Presentation Session 2: Themes 2 & 3</i>	
PROGRESS REPORT	1:15 pm	Al Leon-Garcia / Hadi Bannazadeh - Theme 5, Univ. of Toronto	
	1:30 pm	Marin Litoiu - Theme 5, York University	
Theme 5: SAVI Testbed	1:40 pm	Raouf Boutaba - Theme 5, University of Waterloo	
	1:50 pm	Sudhaker Ganti - Theme 5, University of Victoria	
	2:00 pm	Paul Chow - Theme 5, University of Toronto	
	2:10 pm	<i>Questions and Answers Period</i>	
	2:20 pm	<i>Coffee Break and Presentation Session 3: Themes 4 & 5</i>	
	MEETING	3:15 pm	SAVI Research Program Committee (Bahen Bldg., Room BA7180)

Location: Bahen Building, Room 1190, 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



“SMART APPLICATIONS ON VIRTUAL INFRASTRUCTURE”

2016 Annual General Meeting Agenda: Thursday, July 7, 2016 - 8:30 am to 2:00 pm

8:30 am	<i>Broadband Enabling the Smart City</i> - Ravinder Shergill , Principal Architect, Telus Corporation
9:00 am	<i>Virtualization of Internet of Things (IoT)</i> - Roch Glitho , Associate Profesor, Université Concordia University
9:30 am	Panel Discussion: <i>Transition to Software-Based Infrastructure</i> - David Mann , (Panel Moderator), SAVI Scientific Advisor/Industry Liaison and President, Ayshireton Consulting Inc.; Ravinder Shergill , Principal Technology Architect, Telus Corporation; Chris Bachalo , Chief Techology Officer Canada, Juniper Networks; Jerry Glowka , Vice-President, SynerSolutions; Raouf Boutaba , Professor, University of Waterloo
10:30 am	Break
11:00 am	<i>Smarter Cities: Applications and Networks</i> - Richard McDonald and Michael J. Martin , IBM Canada Ltd.
11:30 am	<i>Software-Defined Fabrics for IoT at Scale</i> - Al Leon-Garcia , Professor and SAVI Scientific Director, University of Toronto
12:00 pm	Panel Discussion: <i>Testbed and Research Networks</i> - Rick McGeer , (Panel Moderator), Chief Scientist, US Ignite and Principal Investigator, Communications Design Group of SAP Labs America; Ted Longley , Director of Engineering & Network Operations, ORION; Mark Wolff , Chief Technology Officer, CANARIE; Mohamed Cheriet , Professor, École de technologie supérieur; Rod Wilson , Senior Director of External Research Programs, Ciena
12:55 pm	<i>Closing of AGM</i> - Bill Hutchison , Hutchison Management International
1:00 pm	Lunch and Demonstrations
2:00 pm	SAVI Board of Directors Meeting (Bahen Building, Room BA7180)

Location: Bahen Building, Room 1190, 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

["Delay-Optimized Video Traffic Routing in Software-Defined Inter-Datacenter Networks"](#) Yinan Liu, Di Niu, **Baochun Li**, (University of Toronto)

["Zebra: Demand-aware Erasure Coding for Distributed Storage Systems"](#) Jun Li and **Baochun Li**, (University of Toronto)

["Bottleneck Detection for Wide Area Data Analytics on the SAVI Testbed"](#) Hao Wang and **Baochun Li**, (University of Toronto)

["Planning Ahead for Future Jobs in Wide Area Data Analytics"](#) Siqi Ji and **Baochun Li**, (University of Toronto)

["Optimizing Coflow Completion Times with Utility Max-Min Fairness"](#) Li Chen, Wei Cui, Bo Li, **Baochun Li**, (University of Toronto)

["Virtual Network Function Description, Publication and Discovery"](#) Sandya Shanmugasundaram, Sami Yangui, **Roch Glitho**, (Université Concordia University)

["A Platform as a Service for Content Delivery Networks"](#) Aida Rangy, Sami Yangui, **Roch Glitho**, (Université Concordia University)

["Real-time Traffic-based Routing, based on Open Data and Open-Source Software"](#) Diego Serrano, Teresa Baldassarre, **Eleni Stroulia**, (University of Alberta)

["Kaleidoscope: A Cloud-Based Platform for Real-Time Video-based Interaction"](#) Hu Zhang, Diego Serrano, **Eleni Stroulia**, (University of Alberta)

["PriMoR-sky: A Runtime Models Processing Infrastructure for Cyber Physical Applications in the Cloud"](#) Lorena Castaneda, Norha Villegas, **Hausi Müller**, (University of Victoria)

["Green Data Centers: Smart Applications and Energy Metrics"](#) Andreas Bergen, **Hausi Müller**, **Sudhakar Ganti**, (University of Victoria)

["CAVA: Context Aware Video Analytics on the SAVI Network"](#) Nina Taherimaksousi, **Hausi Müller**, (University of Victoria)

POSTER & DEMONSTRATION

["Online Group Chat with Bellini on the SAVI Testbed"](#) Shuhao Liu, Yinan Liu, **Baochun Li**, (University of Toronto)

["Cross-Region Data Delivery Service on the SAVI Testbed"](#) Shuhao Liu, Yinan Liu, **Baochun Li**, (University of Toronto)



POSTER & DEMONSTRATION TITLE ROSTER



POSTER & DEMONSTRATION: Theme 1 Cont'd

["An NFV Based Architecture for On-the-fly Provisioning of Value Added Video Services in Content Delivery Networks"](#) Elaheh Narjes Taghig Jahromi, Sandhya Shanmugasundaram, Aida Rangy, Sami Yangui, Mohammad Ali Salahuddin, **Roch Glitho**, (Université Concordia University)

["A Cloud Platform-as-a-Service for Multimedia Conferencing Service Provisioning"](#) Ahmad Ferdous Bin Alam, Abbas Soltanian, Sami Yangui, Mohammad Ali Salahuddin, Halima Elbiaze, **Roch Glitho**, (Université Concordia University)

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

THEME 2 - RESOURCE CONTROL & MANAGEMENT

POSTERS

["SpearDB: Database Management at the Network Edge"](#) Cătălin Avram, Bernard Wong, **Ken Salem**, (University of Waterloo)

["Modeling Performance Impact of Middleware in Cloud Systems"](#) Adnan Faisal, Dorina Petriu, **Murray Woodside**, (Carleton University)

["LEGIS on Scale with Docker Containers"](#) Marios Fokaefs, Cornel Barna, Rodrigo Veleda, **Marin Litoiu**, (York University)

["On Efficiency and Scalability of Software-Defined Infrastructure for Adaptive Applications"](#) Nasim Beigi-Mohammadi, Hamzeh Khazaei, Mark Shtern, Cornel Barna, **Marin Litoiu**, (York University)

["An Adaptive Pricing Model for Scalable Cloud Applications"](#) Marios Fokaefs, Cornel Barna, **Marin Litoiu**, (York University)

["Provisioning Performance of Cloud Microservices"](#) Hamzeh Khazaei, Cornel Barna, Nasim Beigi-Mohammadi, **Marin Litoiu**, (York University)

["Overload Management for Busy Multi-tier Web Applications using Software Defined Networking"](#) Nasim Beigi-Mohammadi, Hamzeh Khazaei, Mark Shtern, Cornel Barna, **Marin Litoiu**, (York University)

POSTER & DEMONSTRATION

["Bitnobi - A Big Data Platform for Distributed Clouds"](#) Mark Shtern, Hassan Jaferi, Ethan Nguyen, Justin Cuaresma, **Marin Litoiu**, (York University)

["Hogna as a DevOps Framework for Containerized Cloud Applications"](#) Cornel Barna, Marios Fokaefs, Hamzeh Khazaei, **Marin Litoiu**, (York University)



POSTER & DEMONSTRATION TITLE ROSTER



DEMONSTRATIONS

["Real-time Detection and Identification of Events on Ontario Highways"](#) Hamzeh Khazaei, Rodrigo Veleza, **Marin Litoiu**, (York University)

THEME 3 - SMART CONVERGED EDGE

POSTERS

["Benchmarking and Evaluating OpenFlow Controllers"](#) Mohamad Darianian and **Carey Williamson**, (University of Calgary)

["Towards a Scalable Monitoring Infrastructure Using SDN"](#) **Carey Williamson** and **Majid Ghaderi**, (University of Calgary)

["Evaluating RTE and NC as Network Services"](#) Feifei Shi and **Carey Williamson**, (University of Calgary)

["Distributed Service Function Chaining"](#) Milad Ghaznavi, Nashid Shahriar, Shahin Kamali, Reaz Ahmed, **Raouf Boutaba**, (University of Waterloo)

["Fabric Embedding with Link Failure Protection Guarantees"](#) Aimal Khan, Nashid Shahriar, Milad Ghaznavi, Shihabur Rahman Chowdhury, Reaz Ahmed, **Raouf Boutaba**, (University of Waterloo)

["Connectivity-aware Virtual Network Embedding"](#) Nashid Shahriar, Reaz Ahmed, Shihabur Rahman Chowdhury, Md. Mashrur Alam Khan, **Raouf Boutaba**, (University of Waterloo)

["On Orchestrating Virtual Network Functions"](#) Md. Faizul Bari, Shihabur Rahman Chowdhury, Reaz Ahmed, **Raouf Boutaba**, (University of Waterloo)

["Revolutionizing Home Firewalls with FleXam"](#) Sajad Shirali-Shahreza and **Yashar Ganjali**, (University of Toronto)

["Dependence of TCP Performance on Congestion Control Parameters"](#) Shiva Ketabi, Monia Ghobadi, Soheil Hassas-Yeganeh, **Yashar Ganjali**, (University of Toronto)

DEMONSTRATIONS

["NF.IO: Network Function Discovery and Orchestration"](#) Md. Faizul Bari, Shihabur Rahman Chowdhury, Reaz Ahmed, **Raouf Boutaba**, (University of Waterloo)



POSTER & DEMONSTRATION TITLE ROSTER



THEME 4 - INTEGRATED WIRELESS/OPTICAL ACCESS

POSTERS

["A New MAC Protocol Design to Support Resource Management in Virtual Wireless Networks"](#) Daniel Tweed, Alfred Kenny, Nhat-Quang Dao, Quang-Dung Ho, **Tho Le-Ngoc**, (McGill University)

["Inter-Access Point Coordination for Efficient Network Resource Management"](#) Nhat-Quang Dao, Daniel Tweed, Alfred Kenny, Quang-Dung Ho, **Tho Le-Ngoc**, (McGill University)

["SiP modulators for Radio Over Fiber"](#) Jiachaun Lin, Bahareh Sherafati, Hassan Sepehrian, Wei Shi, **Leslie A. Rusch**, (Université Laval)

["Constellation-Induced SNR Gain in Short-Reach Optical OFDM"](#) Mingyang Lyu, Siamak Amiralizadeh, **Leslie A. Rusch**, (Université Laval)

["Closing the Loop: Data and RoF Symbiosis in Uplink and Downlink"](#) Zhihui Cao, Jiachuan Lin, **Leslie A. Rusch**, (Université Laval)

POSTER & DEMONSTRATION

["Bringing Control of Radio Resources to Virtualized Networks: The Implementation of a Programmable Medium Access Protocol"](#) Alfred Kenny, Nhat-Quang Dao, Daniel Tweed, Quang-Dung Ho, **Tho Le-Ngoc**, (McGill University)

THEME 5 - APPLICATION PLATFORM TESTBED

POSTERS

["Scalable Reconfigurable FPGA Clusters in the Cloud"](#) Naif Tarafdar, Erik Fukuda, Jiayuan Chen, **Paul Chow**, (University of Toronto)

["Organizing FPGAs in the Cloud for Distributed Applications"](#) Naif Tarafdar, Eric Fukuda, **Paul Chow**, (University of Toronto)

["Cloud Resource Monitoring for Facilitating Administration"](#) Sumit Kadyan and **Sudhakar Ganti**, (University of Victoria)

["Autonomous Power Consumption Scheduling for Smart Homes with Local Power Generation and Storage Facilities"](#) Mohammad Hossein Yaghmaee Moghaddam, Morteza Moghaddassian, **Alberto Leon-Garcia**, (University of Toronto)

["Graph-based Diagnosis in Software-Defined Infrastructure"](#) Joseph Wahba, Hazem Soliman, Hadi Bannazadeh, **Alberto Leon-Garcia**, (University of Toronto)



POSTER & DEMONSTRATION TITLE ROSTER



POSTERS: Theme 5 Cont'd

["Enhanced Real Time Content Delivery using vCPE and NFV Service Chaining"](#) Pouya Yasrebi, Hadi Bannazadeh, **Alberto Leon-Garcia**, (University of Toronto)

["OpenFlow Conflict Detection and Authorization in Multi-Tenant Clouds"](#) Byungchul Park, Thomas Lin, Hadi Bannazadeh, **Alberto Leon-Garcia**, (University of Toronto)

POSTER & DEMONSTRATION

["End-to-End Peering across SAVI and GENI Testbeds using SDI-based Exchange Points"](#) Saeed Arezoumand, Hadi Bannazadeh, **Alberto Leon-Garcia**, (University of Toronto)

["NFV Service Chaining in SAVI Testbed"](#) Saeed Arezoumand, Thomas Lin, Hadi Bannazadeh, **Alberto Leon-Garcia**, (University of Toronto)

["Extending Software-defined Infrastructure to Legacy Systems"](#) Spandan Bemby, Hadi Bannazadeh, **Alberto Leon-Garcia**, (University of Toronto)

Note: Refer to following page for schedule of SAVI Posters and Demonstrations



POSTER & DEMONSTRATION SCHEDULE



WEDNESDAY, JULY 6 - 10:10 AM - SESSION 1 THEME 1: POSTERS

STN.NO. POSTER TITLE

P1	<u>"Delay-Optimized Video Traffic Routing in Software-Defined Inter-Datacenter Networks"</u> Yinan Liu, Di Niu, Baochun Li , (University of Toronto)
P2	<u>"Zebra: Demand-aware Erasure Coding for Distributed Storage Systems"</u> Jun Li and Baochun Li , (University of Toronto)
P3	<u>"Bottleneck Detection for Wide Area Data Analytics on the SAVI Testbed"</u> Hao Wang and Baochun Li , (University of Toronto)
P4	<u>"Planning Ahead for Future Jobs in Wide Area Data Analytics"</u> Siqi Ji and Baochun Li , (University of Toronto)
P5	<u>"Optimizing Coflow Completion Times with Utility Max-Min Fairness"</u> Li Chen, Wei Cui, Bo Li, Baochun Li , (University of Toronto)
P6	<u>"Virtual Network Function Description, Publication and Discovery"</u> Sandya Shanmugasundaram, Sami Yangui, Roch Glitho , (Université Concordia University)
P7	<u>"A Platform as a Service for Content Delivery Networks"</u> Aida Rangy, Sami Yangui, Roch Glitho , (Université Concordia University)
P8	<u>"Real-time Traffic-based Routing, based on Open Data and Open-Source Software"</u> Diego Serrano, Teresa Baldassarre, Eleni Stroulia , (University of Alberta)
P9	<u>"Kaleidoscope: A Cloud-Based Platform for Real-Time Video-based Interaction"</u> Hu Zhang, Diego Serrano, Eleni Stroulia , (University of Alberta)
P10	<u>"PrIMoR-sky: A Runtime Models Processing Infrastructure for Cyber Physical Applications in the Cloud"</u> Lorena Castaneda, Norha Villegas, Hausi Müller , (University of Victoria)
P11	<u>"Green Data Centers: Smart Applications and Energy Metrics"</u> Andreas Bergen, Hausi Müller , Sudhakar Ganti , (University of Victoria)
P12	<u>"CAVA: Context Aware Video Analytics on the SAVI Network"</u> Nina Taherimakhsousi, Hausi Müller , (University of Victoria)

WEDNESDAY, JULY 6 - 10:10 AM - SESSION 1 THEME 1: POSTERS & DEMONSTRATIONS

STN.NO. POSTER & DEMONSTRATION TITLE

PD1	<u>"Online Group Chat with Bellini on the SAVI Testbed"</u> Shuhao Liu, Yinan Liu, Baochun Li , (University of Toronto)
-----	--



POSTER & DEMONSTRATION SCHEDULE



STN.NO. POSTER & DEMONSTRATION TITLE - 10:10 am, Session 1, Cont'd

- PD1 ["Cross-Region Data Delivery Service on the SAVI Testbed"](#) Shuhao Liu, Yinan Liu, Baochun Li, (University of Toronto)
- PD2 ["An NFV Based Architecture for On-the-fly Provisioning of Value Added Video Services in Content Delivery Networks"](#) Elaheh Narjes Taghigh Jahromi, Sandhya Shanmugasundaram, Aida Rangy, Sami Yangui, Mohammad Ali Salahuddin, **Roch Glitho**, (Université Concordia University)
- PD3 ["A Cloud Platform-as-a-Service for Multimedia Conferencing Service Provisioning"](#) Ahmad Ferdous Bin Alam, Abbas Soltanian, Sami Yangui, Mohammad Ali Salahuddin, Halima Elbiaze, **Roch Glitho**, (Université Concordia University)
- PD4 ["Kaleidoscope --- Applications for Multi-Tier Cloud, Smart Edge/Core and Adaptive Resource Support"](#) Ronald Desmarais, Andreas Bergen, **Hausi Müller**, **Sudhakar Ganti**, (University of Victoria)

WEDNESDAY, JULY 6 - 12:30 PM - SESSION 2 THEME 2 & 3: POSTERS

STN.NO. POSTER TITLE

- P1 ["SpearDB: Database Management at the Network Edge"](#) Cătălin Avram, Bernard Wong, **Ken Salem**, (University of Waterloo)
- P2 ["Modeling Performance Impact of Middleware in Cloud Systems"](#) Adnan Faisal, Dorina Petriu, **Murray Woodside**, (Carleton University)
- P3 ["LEGIS on Scale with Docker Containers"](#) Marios Fokaefs, Cornel Barna, Rodrigo Veleda, **Marin Litoiu**, (York University)
- P4 ["On Efficiency and Scalability of Software-Defined Infrastructure for Adaptive Applications"](#) Nasim Beigi-Mohammadi, Hamzeh Khazaei, Mark Shtern, Cornel Barna, **Marin Litoiu**, (York University)
- P5 ["An Adaptive Pricing Model for Scalable Cloud Applications"](#) Marios Fokaefs, Cornel Barna, **Marin Litoiu**, (York University)
- P6 ["Provisioning Performance of Cloud Microservices"](#) Hamzeh Khazaei, Cornel Barna, Nasim Beigi-Mohammadi, **Marin Litoiu**, (York University)
- P7 ["Overload Management for Busy Multi-tier Web Applications using Software Defined Networking"](#) Nasim Beigi-Mohammadi, Hamzeh Khazaei, Mark Shtern, Cornel Barna, **Marin Litoiu**, (York University)



POSTER & DEMONSTRATION SCHEDULE



STN.NO. POSTER TITLE - 12:30 pm, Session 2, Cont'd

- P8 ["Benchmarking and Evaluating OpenFlow Controllers"](#) Mohamad Darianian and **Carey Williamson**, (University of Calgary)
- P9 ["Towards a Scalable Monitoring Infrastructure Using SDN"](#) **Carey Williamson** and **Majid Ghaderi**, (University of Calgary)
- P10 ["Evaluating RTE and NC as Network Services"](#) Feifei Shi and **Carey Williamson**, (University of Calgary)
- P11 ["Distributed Service Function Chaining"](#) Milad Ghaznavi, Nashid Shahriar, Shahin Kamali, Reaz Ahmed, **Raouf Boutaba**, (University of Waterloo)
- P12 ["Fabric Embedding with Link Failure Protection Guarantees"](#) Aimal Khan, Nashid Shahriar, Milad Ghaznavi, Shihabur Rahman Chowdhury, Reaz Ahmed, **Raouf Boutaba**, (University of Waterloo)
- P13 ["Connectivity-aware Virtual Network Embedding"](#) Nashid Shahriar, Reaz Ahmed, Shihabur Rahman Chowdhury, Md. Mashrur Alam Khan, **Raouf Boutaba**, (University of Waterloo)
- P14 ["On Orchestrating Virtual Network Functions"](#) Md. Faizul Bari, Shihabur Rahman Chowdhury, Reaz Ahmed, **Raouf Boutaba**, (University of Waterloo)
- P15 ["Revolutionizing Home Firewalls with FlexAm"](#) Sajad Shirali-Shahreza and **Yashar Ganjali**, (University of Toronto)
- P16 ["Dependence of TCP Performance on Congestion Control Parameters"](#) Shiva Ketabi, Monia Ghobadi, Soheil Hassas-Yeganeh, **Yashar Ganjali**, (University of Toronto)

WEDNESDAY, JULY 6 - 12:30 PM - SESSION 2 THEME 2 & 3: POSTERS & DEMONSTRATIONS

STN.NO. POSTER & DEMONSTRATION TITLE

- PD1 ["Bitnobi - A Big Data Platform for Distributed Clouds"](#) Mark Shtern, Hassan Jaferi, Ethan Nguyen, Justin Cuaresma, **Marin Litoiu**, (York University)
- PD2 ["Hogna as a DevOps Framework for Containerized Cloud Applications"](#) Cornel Barna, Marios Fokaefs, Hamzeh Khazaei, **Marin Litoiu**, (York University)
-

WEDNESDAY, JULY 6 - 12:30 PM - SESSION 2 THEME 2 & 3: DEMONSTRATIONS

STN.NO. DEMONSTRATION TITLE

- D1 ["Real-time Detection and Identification of Events on Ontario Highways"](#) Hamzeh Khazaei, Rodrigo Veleda, **Marin Litoiu**, (York University)
- D2 ["NF.IO: Network Function Discovery and Orchestration"](#) Md. Faizul Bari, Shihabur Rahman Chowdhury, Reaz Ahmed, **Raouf Boutaba**, (University of Waterloo)

WEDNESDAY, JULY 6 - 2:20 PM - SESSION 3 THEMES 4 & 5: POSTERS

STN.NO. POSTER TITLE

- P1 ["A New MAC Protocol Design to Support Resource Management in Virtual Wireless Networks"](#) Daniel Tweed, Alfred Kenny, Nhat-Quang Dao, Quang-Dung Ho, **Tho Le-Ngoc**, (McGill University)
- P2 ["Inter-Access Point Coordination for Efficient Network Resource Management"](#) Nhat-Quang Dao, Daniel Tweed, Alfred Kenny, Quang-Dung Ho, **Tho Le-Ngoc**, (McGill University)
- P3 ["SiP modulators for Radio Over Fiber"](#) Jiachaun Lin, Bahareh Sherafati, Hassan Sepehrian, Wei Shi, **Leslie A. Rusch**, (Université Laval)
- P4 ["Constellation-Induced SNR Gain in Short-Reach Optical OFDM"](#) Mingyang Lyu, Siamak Amiralizadeh, **Leslie A. Rusch**, (Université Laval)
- P5 ["Closing the Loop: Data and RoF Symbiosis in Uplink and Downlink"](#) Zhihui Cao, Jiachuan Lin, **Leslie A. Rusch**, (Université Laval)
- P6 ["Scalable Reconfigurable FPGA Clusters in the Cloud"](#) Naif Tarafdar, Erik Fukuda, Jiayuan Chen, **Paul Chow**, (University of Toronto)
- P7 ["Organizing FPGAs in the Cloud for Distributed Applications"](#) Naif Tarafdar, Eric Fukuda, **Paul Chow**, (University of Toronto)
- P8 ["Cloud Resource Monitoring for Facilitating Administration"](#) Sumit Kadyan and **Sudhakar Ganti**, (University of Victoria)
- P9 ["Autonomous Power Consumption Scheduling for Smart Homes with Local Power Generation and Storage Facilities"](#) Mohammad Hossein Yaghmaee Moghaddam, Morteza Moghaddassian, **Alberto Leon-Garcia**, (University of Toronto)
- P10 ["Graph-based Diagnosis in Software-Defined Infrastructure"](#) Joseph Wahba, Hazem Soliman, Hadi Bannazadeh, **Alberto Leon-Garcia**, (University of Toronto)



POSTER & DEMONSTRATION SCHEDULE



STN.NO. POSTER TITLE 2:20 pm, Session 3, Cont'd

- P11 ["Enhanced Real Time Content Delivery using vCPE and NFV Service Chaining"](#) Pouya Yasrebi, Hadi Bannazadeh, **Alberto Leon-Garcia**, (University of Toronto)
- P12 ["OpenFlow Conflict Detection and Authorization in Multi-Tenant Clouds"](#) Byungchul Park, Thomas Lin, Hadi Bannazadeh, **Alberto Leon-Garcia**, (University of Toronto)

WEDNESDAY, JULY 6 - 2:20 PM - SESSION 3 THEMES 4 & 5: POSTERS & DEMONSTRATIONS

STN.NO. POSTER & DEMONSTRATION TITLE

- PD1 ["Bringing Control of Radio Resources to Virtualized Networks: The Implementation of a Programmable Medium Access Protocol"](#) Alfred Kenny, Nhat-Quang Dao, Daniel Tweed, Quang-Dung Ho, **Tho Le-Ngoc**, (McGill University)
- PD2 ["End-to-End Peering across SAVI and GENI Testbeds using SDI-based Exchange Points"](#) Saeed Arezoumand, Hadi Bannazadeh, **Alberto Leon-Garcia**, (University of Toronto)
- PD3 ["NFV Service Chaining in SAVI Testbed"](#) Saeed Arezoumand, Thomas Lin, Hadi Bannazadeh, **Alberto Leon-Garcia**, (University of Toronto)
- PD4 ["Extending Software-defined Infrastructure to Legacy Systems"](#) Spandan Bemby, Hadi Bannazadeh, **Alberto Leon-Garcia**, (University of Toronto)

THURSDAY, JULY 7 - 1:00 PM - SINGLE SESSION ALL DEMONSTRATIONS & POSTER/DEMONSTRATION COMBINATIONS

STN.NO. POSTER & DEMONSTRATION TITLE

- PD1 ["Online Group Chat with Bellini on the SAVI Testbed"](#) Shuhao Liu, Yinan Liu, **Baochun Li**, (University of Toronto)
- PD1 ["Cross-Region Data Delivery Service on the SAVI Testbed"](#) Shuhao Liu, Yinan Liu, **Baochun Li**, (University of Toronto)
- PD2 ["An NFV Based Architecture for On-the-fly Provisioning of Value Added Video Services in Content Delivery Networks"](#) Elaheh Narjes Taghieh Jahromi, Sandhya Shanmugasundaram, Aida Rangy, Sami Yangui, Mohammad Ali Salahuddin, **Roch Glitho**, (Université Concordia University)
- PD3 ["A Cloud Platform-as-a-Service for Multimedia Conferencing Service Provisioning"](#) Ahmad Ferdous Bin Alam, Abbas Soltanian, Sami Yangui, Mohammad Ali Salahuddin, Halima Elbiaze, **Roch Glitho**, (Université Concordia University)



POSTER & DEMONSTRATION SCHEDULE



STN.NO. POSTER & DEMONSTRATION TITLE, 1:00 pm Cont'd

- PD4 ["Kaleidoscope --- Applications for Multi-Tier Cloud, Smart Edge/Core and Adaptive Resource Support"](#) Ronald Desmarais, Andreas Bergen, Hausi Müller, Sudhakar Ganti, (University of Victoria)
- PD5 ["Bitnobi - A Big Data Platform for Distributed Clouds"](#) Mark Shtern, Hassan Jaferi, Ethan Nguyen, Justin Cuaresma, **Marin Litoiu**, (York University)
- PD6 ["Hogna as a DevOps Framework for Containerized Cloud Applications"](#) Cornel Barna, Marios Fokaefs, Hamzeh Khazaei, **Marin Litoiu**, (York University)
- PD7 ["Bringing Control of Radio Resources to Virtualized Networks: The Implementation of a Programmable Medium Access Protocol"](#) Alfred Kenny, Nhat-Quang Dao, Daniel Tweed, Quang-Dung Ho, **Tho Le-Ngoc**, (McGill University)
- PD8 ["End-to-End Peering across SAVI and GENI Testbeds using SDI-based Exchange Points"](#) Saeed Arezoumand, Hadi Bannazadeh, **Alberto Leon-Garcia**, (University of Toronto)
- PD9 ["NFV Service Chaining in SAVI Testbed"](#) Saeed Arezoumand, Thomas Lin, Hadi Bannazadeh, **Alberto Leon-Garcia**, (University of Toronto)
- PD10 ["Extending Software-defined Infrastructure to Legacy Systems"](#) Spandan Bemby, Hadi Bannazadeh, **Alberto Leon-Garcia**, (University of Toronto)

STN.NO. DEMONSTRATION TITLE

- D1 ["Real-time Detection and Identification of Events on Ontario Highways"](#) Hamzeh Khazaei, Rodrigo Veleda, **Marin Litoiu**, (York University)
- D2 ["NF.IO: Network Function Discovery and Orchestration"](#) Md. Faizul Bari, Shihabur Rahman Chowdhury, Reaz Ahmed, **Raouf Boutaba**, (University of Waterloo)

Note: ["Online Group Chat with Bellini on the SAVI Testbed"](#) and ["Cross-Region Data Delivery Service on the SAVI Testbed"](#) will be presented together

GUEST SPEAKER BIO

Mr. Ravinder Shergill

Principal Architect
Telus Corporation



Ravinder Shergill is the Principal Architect responsible for Strategy, Architecture and Governance for IP infrastructures at TELUS.

With his deep, 25+ years, experience in networking, Ravinder helps set the strategy for next-generation architectures that promote industry ecosystem adoption, builds cross-segment synergies between different 'tenants' of the network.

He has served in a variety of lead roles at TELUS, including Converged Core Architect, Converged Edge Architect, IPv6 enablement, CO Rationalization, defining the next generation topology in an all-Fiber and all-IP 'application centric' infrastructure.

As the Chief Architect, Ravinder is presently the prime for SDN Architecture and Strategy for TELUS. Having lead a highly converged architecture from the core of the network to the edge over the past decade+, now he's rationalizing convergence in the access sections of the network that leverages Cloud, Virtualization, and Softwarization, paving the way for massive connectivity for emerging technologies.

GUEST SPEAKER BIO

Dr. Roch Glitho

Associate Professor of Networking & Telecommunications
Université Concordia University



Dr. Roch Glitho is an associate professor of networking and telecommunications at the Concordia Institute of Information Systems Engineering (CIISE), Université Concordia University, Montreal, Canada where he holds a Canada Research Chair in End-User Service Engineering for Communication Networks.

Dr. Glitho [SM] (<http://users.encs.concordia.ca/~glitho/>) holds a Ph.D. (Tekn. Dr.) in teleinformatics (Royal Institute of Technology, Stockholm, Sweden) and M.Sc. degrees in business economics (University of Grenoble, France), pure mathematics (University Geneva, Switzerland), and computer science (University of Geneva).

In the past, he has worked in industry for almost a quarter of a century and has held several senior technical positions at LM Ericsson in Sweden and Canada (e.g. expert, principal engineer, senior specialist). Dr. Glitho has also served as IEEE Communications Society distinguished lecturer, Editor-In-Chief of IEEE Communications Magazine and Editor-In-Chief of IEEE Communications Surveys & Tutorials.

Related site:

<http://users.encs.concordia.ca/~tse/people.php>

GUEST SPEAKER BIO

Mr. David Mann

SAVI Scientific Advisor & Industry Liaison
Ayrshireton Consulting Inc. (President)



Mr. David Mann retired in 2001 from Nortel Networks. He was Vice President of Emerging Business Technology Investments. David has lived in seven countries during his career: Scotland, England, Zambia, Sweden, Belgium, Spain and Canada.

Mr. Mann has served on the board of OCRI from 1995-2000 (now Invest Ottawa, <http://investottawa.ca>). He is an honorary member and a past Vice-President of the Canadian Association for the Advancement of Science, <http://caas-acascience.org>.

Currently, David serves on several Boards in Canada:

- President of FM-CFS Canada, www.fm-cfs.ca;
- Advisory Board Chair of InBay Technologies Inc., www.inbaytech.com;
- David also has his own High Tech and business consultancy company (Ayrshireton Consulting Inc.) in Canada, www.ayrshireton.ca;
- Chairman, IET Ottawa Local Network, <https://communities.theiet.org/communities/home/251>

David is the Scientific Advisor and Industry Liaison to the Canadian national research project “Smart Applications on Virtual Infrastructure (SAVI)”.

David was educated in Scotland finishing his post-secondary education at Bell College of Technology. He also received parallel post-secondary education with British Telecom (BT).

David is a Fellow of the Institute of Engineering Technology (IET), a Chartered Engineer (UK) and a Senior Member of the Institute of Electrical and Electronics Engineers (IEEE).

GUEST SPEAKER BIO

Mr. Chris Bachalo

Chief Technology Officer
Juniper Networks Canada



Chris Bachalo is currently the Chief Technical Officer for Juniper Networks Canada. Mr Bachalo has accumulated 25 years of experience in the Canadian telecom industry.

Chris joined Juniper in 2010 from Alcatel-Lucent (and formerly Newbridge Networks) where he was VP of Systems Engineering for 17 years. Chris is focused on building telecommunications solutions for Canadian service providers and large enterprises.

He is also heavily engaged in Canadian academic research and has participated in multiple technology advisory boards including: University of Toronto Masters of Engineering in Telecommunications Advisory Board, NodeLogic Networks Technical Advisory Board and currently serves on Algonquin College Photonics and Laser Technology Advisory Board, BRIC/University of Regina (Bridging Research and Interoperability Center) Industry Advisory Board, SAVI/University of Toronto (Smart Applications on Virtual Infrastructure) Industry Advisory Board and Research Program Committee, CENGN (Center of Excellence for Next Generation Networking) Industry Advisory Board, Invest Ottawa ICT Industry Sector Advisor, BCNet Industry Advisor, Nuvollo Networks industry advisor and is cofounder/technology advisor/co-author of primary patent for Emermotion Inc, a Canadian cleantech company.

Chris has a BAsC in Electrical Engineering from the University of Windsor and studied at the University of Toronto towards a MEng degree in Electrical Engineering.

GUEST SPEAKER BIO

Mr. Jerry Glowka

Vice-President, Emerging Technologies
SynerSolutions



Jerry has been involved with Information Technology for over 25 years. His expertise is in Big Data Analytics, High Performance Computing and in Next Generation Security. Jerry has been a member of SAVI project for the past five years.

Some of his latest accomplishments have been building Canada's largest Big Data Clusters as well as securing government facilities around the world. He is known for bringing new technologies to light to help Government Agencies with their fight against bad actors and adversaries on their secure networks.

He is presently involved in secure Internet of Things (IoT) project with the Federal government.

For more on information on Jerry and SynerSolutions, please visit:

<http://www.synersolutions.com/en/>

GUEST SPEAKER BIO

Dr. Raouf Boutaba

Professor, FIEEE, FIEC, FCAE
University of Waterloo



Dr. Raouf Boutaba received the M.Sc. and Ph.D. degrees in computer science from the University Pierre & Marie Curie, Paris, in 1990 and 1994, respectively. He is currently a professor of computer science at the University of Waterloo. His primary research interests are in network and service management.

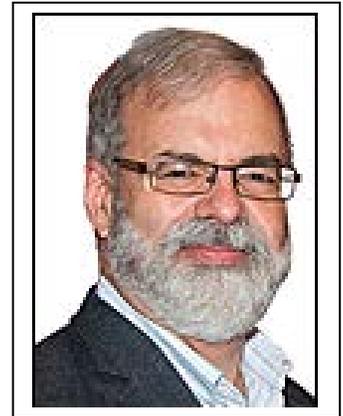
Dr. Boutaba has published extensively in these areas and received several journal and conference Best Paper Awards including the IEEE 2008 Fred W. Ellersick Prize Paper Award. He received several other recognitions such as the Premier's Research Excellence Award, the Nortel Research Excellence Award, the Nortel Excellence Award in Technology Transfer, fellowships of the Faculty of Mathematics and the David R. Cheriton School of Computer Science and several outstanding performance awards at the University of Waterloo. He has also received the IEEE Communications Society Hal Sobol Award and the IFIP Silver Core in 2007, the IEEE Communications Society Joe LociCero and the Dan Stokesbury awards in 2009, the IEEE Communications Society Salah Aidarous award in 2012, and the IEEE Canada McNaughton Gold Medal in 2014.

Dr. Boutaba is the founding editor in chief of the IEEE Transactions on Network and Service Management (2007 - 2010), on the editorial boards of several journals, and served as general or technical program chair of a number of major IEEE conferences. He served as a distinguished lecturer for the IEEE Computer and Communications Societies. He is fellow of the IEEE, the Engineering Institute of Canada, and the Canadian Academy of Engineering.

GUEST SPEAKER BIO

Mr. Richard McDonald

Distinguished IT Architect & Technical Executive
IBM Canada Ltd.



Richard McDonald is an IBM Distinguished Engineer with 34 years of IT experience. Mr. McDonald has spent the first half of his career in the IBM Toronto Laboratory developing products for the company's clients.

Currently, Richard is CTO for IBM North America Technical Sales. He designs innovative solutions for some of IBM's largest clients to help them address business and IT challenges and opportunities.

Over the years, he has developed methods for helping clients identify innovation opportunities within IT. His technical expertise is focused on telco, mobile computing and Internet of Things solutions. Richard sits on the Scientific Advisory Council for the Southern Ontario Smarter Computing Innovation Platform (SOSCIP) - a consortium of universities, government and IBM to advance research in key areas, <http://soscip.org>.

During his spare time, he motorcycles and plays in a jazz band. Each spring Richard participates as a judge at FIRST Robotics and FIRST Lego League events.

GUEST SPEAKER BIO

Mr. Michael J. Martin

MA, MBA, MEd, GDM, SCPM, PMP

Senior Executive Consultant

IBM Canada Ltd.



Michael Martin is a Senior Executive Consultant in the IBM Global Technical Services for Network Services focused on smarter cities, energy & utilities, mining, oil & gas, and transportation solutions based in Toronto, Canada.

Mr. Martin consults primarily on complex RF systems and corporate strategy using systems for machine-to-machine and Internet of Things (IoT) communication networks. Michael's work involves technologies such as WiMAX, LTE, microwave, satellite, optical fibre, two-way radio, and coaxial cabling used for M2M and IoT communications. He has helped clients in 25 countries around the globe in the past eight years.

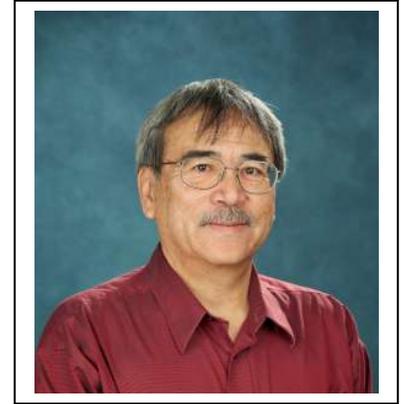
Michael is a dedicated life-long learner currently holding three Masters Degrees. He holds advanced graduate degrees in education, business, and communications, all with emphasis on communication technology, as well as three diplomas and five certificates in a variety of subjects focused on communication and media technologies.

Michael has presented at over 125 conferences and authored more than 100 articles, two book chapters, and several business and technical white papers. He has served for more than 25 years on five college and university boards, for 18 years on a non-profit industry board, as well as on publicly traded company boards. He is currently a director for TeraGo Networks (TSX:TGO).

GUEST SPEAKER BIO

Dr. Alberto Leon-Garcia

Professor & SAVI Project Scientific Director
University of Toronto



Professor Alberto Leon-Garcia is Distinguished Professor in Electrical and Computer Engineering at the University of Toronto. He is recognized as an innovator in networking research and education.

Professor Leon-Garcia is a Fellow of the Institute of Electronics and Electrical Engineering (IEEE) "For contributions to multiplexing and switching of integrated services traffic". He is also a Fellow of the Engineering Institute of Canada (EIC) and the American Association for the Advancement of Science (AAAS). He has received the 2006 Thomas Eadie Medal from the Royal Society of Canada and the 2010 IEEE Canada A. G. L. McNaughton Gold Medal for his contributions to the area of communications. He holds several patents and has published research extensively in the areas of switch architecture and traffic management.

Professor Leon-Garcia is author of the leading textbooks: "Probability and Random Processes for Electrical Engineering" and "Communication Networks: Fundamental Concepts and Key Architecture".

He is Scientific Director of the Natural Sciences and Engineering Research Council of Canada (NSERC) Strategic Network for "Smart Applications on Virtual Infrastructures", www.savinetwork.ca. Professor Leon-Garcia also is the Principal Investigator of an Ontario Research Fund: Research Excellence (ORF/RE) project on "Connected Vehicles and Smart Transportation", www.cvst.ca.

GUEST SPEAKER BIO

Dr. Hadi Bannazadeh

SAVI Project Chief Testbed Architect
University of Toronto



Dr. Bannazadeh earned his PhD from the University of Toronto's Department of Electrical & Computer Engineering. After his graduation, Hadi 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.

GUEST SPEAKER BIO

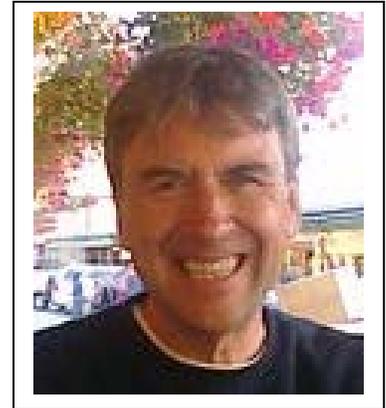
Mr. Rick McGeer

Chief Scientist,

US Ignite

Principal Investigator,

SAP Labs America, Communications Design Group



Rick McGeer is currently Chief Scientist at US Ignite and a Principal Investigator at the Communication Design Group of SAP Labs America.

Previously, he was a Principal Investigator (PI) at the Communication Design Group of SAP Labs America, Distinguished Technologist in HP Enterprise Systems and HP Laboratories, the co-founder and Chief Scientist of Softface, Inc., the co-founder and Research Scientist at Cadence Berkeley Labs, a Research Scientist at UC-Berkeley and an Assistant Professor at the University of British Columbia. He is an Adjunct Professor of Computer Science at the University of Victoria. Rick earned his Ph. D. in Computer Science from UC-Berkeley. He is the author of two books and over 100 refereed papers in technical conferences and journals.

Rick was a co-founder of the PlanetLab consortium and remains a member of the PlanetLab Steering Committee. He was on the original GENI design team and led the InstaGENI initiative, and currently leads the GENI Experiment Engine (Ignite App Engine) project. He is the co-editor of the upcoming book *The GENI Book*, due from Springer-Verlag in July 2016. Rick has been a PI on several DARPA (Defense Advanced Research Projects Agency) programs, notably the Global Mobile (GloMo) program and the Control Plane Program, where he and Jack Brassil led the CHART team to achieve a mean 40x performance improvement on the TCP/IP protocol under military conditions.

To learn more about GENI, please visit: <https://www.geni.net/>.

To learn more about DARPA, please visit: <http://www.darpa.mil/>.

GUEST SPEAKER BIO

Mr. Ted Longley

*Director of Engineering and Network Operations,
Ontario Research and Innovation Optical Network (ORION)*



Ted holds a BA from Sir Wilfred Laurier University. He is an active member of both the International Coach Federation (ICF) as well as the Internet of Things and Cloud Computing (IOTCC).

As a member of the senior management team at ORION, Ted provides strategic leadership for the engineering and network operations department.

Prior to joining ORION in 2015, Ted spent six years with BlackBerry, where he was an early thought leader and champion for Software Defined Networking/Network Programmability, creating and developing a formal, R&D-focused architectural function within BlackBerry global networks. He led the team that brought the re-architecture of the BlackBerry global network to fruition, resulting in significant reductions in routing complexity, better operational manageability, the increased ability to rapidly provision new services, and enhanced security capabilities.

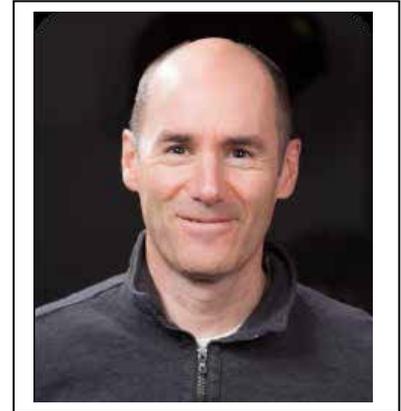
Most recently, Ted was Vice President of Platform Services at ThinkingPhones, where he globalized the multi-tenant cloud-based service offering by opening new data centres in Europe and the US. He also introduced and integrated automation and continuous integration (CI) concepts, which significantly reduced delivery time of new systems and also reduced implementation error rates.

To learn more about ORION, please visit:

<http://www.orion.on.ca>

GUEST SPEAKER BIO

Mr. Mark Wolff,
Chief Technology Officer,
Canadian Network for the Advancement of
Research, Industry and Education (CANARIE)



Mark joined CANARIE in May 2011 as Senior Director, Technology Innovation.

As CTO, Mark's responsibility is to ensure CANARIE continues to drive transformative change in the use of digital infrastructure among Canada's research, education and innovation communities.

Prior to joining CANARIE, Mark has held innovation, development and leadership positions at Newbridge Networks, Alcatel, and Meriton Networks. Mark earned a Bachelor of Commerce and a B.Sc. degree in Electrical Engineering from the University of Saskatchewan. He also holds several patents in telecommunications.

An accomplished guitarist and musician who hails from Saskatoon, Mark now makes his home in Ottawa, where he is an avid snowboarder and mountain biker.

For more information on CANARIE, please visit:

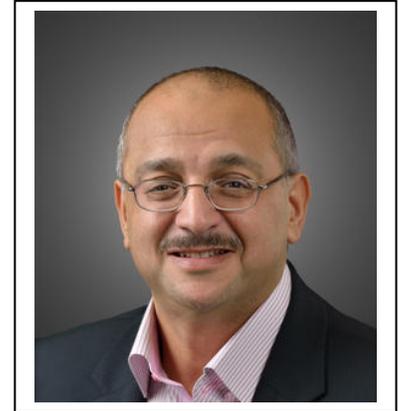
<http://www.canarie.ca/>

GUEST SPEAKER BIO

Dr. Mohamed Cheriet

Professor, Eng., PhD, SMIEEE

École de technologie supérieur (ÉTS)



Professor Cheriet has published more than 300 technical papers in renown international journals and conferences and has delivered more than 30 invited talks.

In addition, he has authored and published 6 books on pattern recognition, document image analysis and understanding, and computer vision. Among them, the book entitled “Character Recognition Systems, a Textbook for Students and Practitioners” is highly acclaimed.

Prof. Cheriet is also recognized for his activities in technical journal editorial writing, organizing and taking part in many conferences. He has contributed to the training of more than 70 high qualified personnel. Dr. Cheriet was awarded the Queen Elizabeth II Diamond Jubilee Medal in light of his significant contributions to knowledge improvement in computational intelligence and mathematical modeling for image processing, created by MITCAS to mark the 60th anniversary of Her Majesty’s accession to the throne. He holds NSERC Canada Research Chair Tier 1 in Sustainable Smart echo-Cloud.

Dr. Cheriet is a senior member of the IEEE and the chapter founder and former chair of IEEE Montreal Computational Intelligent Systems (CIS).

To learn more about Professor Cheriet and his research interests, please visit:

<http://profs.etsmtl.ca/mcheriet/index.html>

GUEST SPEAKER BIO

Mr. Rodney G. Wilson

Sr. Director External Research & Technology Incubation Programs
Ciena Corporation, Ottawa R & D Labs



Mr. Wilson is responsible for Ciena's leadership & interactions with universities and the research community, including national research and education networks.

Employed within Ciena's CTO Group, he also leads Ciena's Technology Incubation Program which seeks to kick-start new development vectors for the company. 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 other roles including director of Broadband Switching, and Optical Ethernet development. Previous roles include University of Toronto as chief network architect for the University's on-line Library business UTLAS.

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

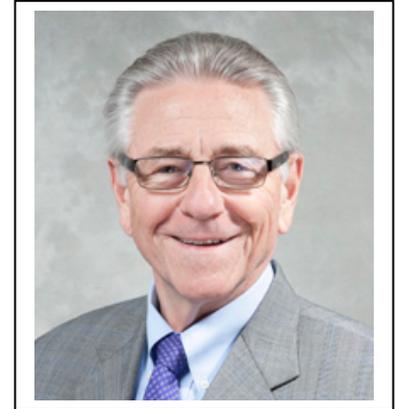
He has served on a number of business and volunteer Boards. He was awarded the Queen Elizabeth II Diamond Jubilee Medal for his service to Canada.

GUEST SPEAKER BIO

Mr. William G. (Bill) Hutchison

F.C.I.P.S., P.Eng.

Hutchison Management International



Bill Hutchison is a business and social entrepreneur, corporate director and international management consultant. He has been Global Chairman, Telecom Industry for Ernst & Young, founder of four successful business start-ups and CEO of three corporate turnarounds.

As a social entrepreneur he has been the founding chair or founding director of ten industry and social consortia and charitable foundations. As technology has evolved Bill Hutchison and his colleagues have been at the forefront in globally managing or advising business and governments on how to achieve the twin goals of transformation and prosperity. Activities have included the creation and restructuring of companies, policy advice to governments on new smart communities and leading the creation of public/private sector consortia to improve national and local competitiveness.

You may download Bill's complete business overview and professional resume here: wghutchison-cv.pdf.

Bill's education, career history, accomplishments and initiatives are also detailed in the Business Overview, Business and Community Initiatives, Awards and Community and Corporate Governance subsections of the Hutchison Management International website:

<http://www.hutchison-management.com>

ABSTRACTS FOR POSTERS & DEMONSTRATIONS

THEME 1 - FUTURE APPLICATIONS

➤ POSTERS:

▪ *Delay-Optimized Video Traffic Routing in Software-Defined Inter-Datacenter Networks*

Yinan Liu, Di Niu, **Baochun Li**, (University of Toronto)

The SAVI testbed provides a wide area infrastructure, which is a perfect platform for many video streaming applications providing geo-distributed services. Through taking advantage of superior connectivity across geographically distributed datacenters presented by the SAVI testbed, the video content is located closer to users and, accordingly, the performance is improved. Concurrently, there are also high volumes of other types of traffic being sent across the inter-datacenter networks, including service replication and data backups, e.g., for storage and email services. It is, therefore, imperative to optimally engineer and schedule inter-datacenter traffic, taking into account the stringent latency requirements of video flows when transmitted along inter-datacenter links shared with other types of traffic.

Hence, in this work, we propose a delay-optimized traffic routing scheme to explicitly differentiate path selection for different flows according to their specific delay-sensitivities. Our solution can effectively respect higher priority flows with lower packet delay, and yield sparse path selection through solving linear programs. In contrast to prior traffic engineering solutions, it does not lead to overly fine-grained traffic splitting, further reducing the packet resequencing overhead and the number of forwarding rules to be installed in each forwarding unit. It is worth mentioning that our sparse traffic routing mechanism is implemented at the application layer based on the software-defined networking (SDN) framework.

▪ *Zebra: Demand-aware Erasure Coding for Distributed Storage Systems*

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

Erasure coding (<http://searchstorage.techtarget.com/definition/erasure-coding>) has been increasingly replacing replication in distributed storage systems, thanks to its lower storage overhead with the same level of failure tolerance. However, with lower storage overhead, the reconstruction overhead of erasure codes can increase significantly as well. Under the ever-changing workload, in which the data access can be highly skewed, it is difficult to achieve a well-defined trade-off between the storage overhead and the reconstruction overhead.

In this work, we propose Zebra, a framework that encodes data with RS codes (https://www.cs.cmu.edu/~guyb/realworld/reedsolomon/reed_solomon_codes.html) into multiple tiers relative to their demand. Given the overall storage overhead and the number of failures to tolerate, Zebra determines the parameters of erasure codes in each tier by solving a geometric programming problem. Based on the demand of data, Zebra can dynamically assign data into the corresponding tiers to minimize the overall reconstruction overhead, and achieve a flexible trade-off between the storage overhead and the reconstruction overhead in multiple tiers, such that hot data can enjoy less overhead of reconstruction and cold data can be stored with lower storage overhead. For the hot data, a better load balance can be achieved with Zebra as well. When demand changes, Zebra can adjust itself accordingly with a marginal amount of network transfer.

▪ *Bottleneck Detection for Wide Area Data Analytics on the SAVI Testbed*

Hao Wang and **Baochun Li**, (University of Toronto)

There have been abundant research work devoted to improving the performance of big data analytics, but few of them take a step further to identify the performance bottlenecks. Existing performance optimization strategies are typically based on priori experience of data analytics applications. e.g., enforcing task placement to save inter-datacenter traffic assumes the network is a dominant bottleneck. However, bottlenecks vary widely from different data analytics applications to different hardware/software configurations.

In this work, we propose a framework that detects performance bottlenecks on-the-flight and makes scheduling decisions based on the actual bottleneck. The geographically distributed computing resources in the SAVI testbed are interconnected by tunneled links in the wide area network, which provides an ideal environment for developing and testing our bottleneck detection framework.

▪ *Planning Ahead for Future Jobs in Wide Area Data Analytics*

Siqi Ji and **Baochun Li**, (University of Toronto)

Big data analytics is widely used for cloud services to process large volumes of data. Since data size is growing exponentially, large organizations are deploying datacenters in a geo-distributed fashion. The centralized approach, which aggregates data in a single datacenter for analytic functions, is no longer efficient. There is a number of research works proposed to perform data analytics in a distributed way that achieves performance benefits. Most big data jobs contain complex shuffle and join, which can cause a large amount of data transfers between datacenters. It has been showed that for recurring jobs, future job characteristics can be predicted with a low error.

In this work, we propose a framework that gives strategies for future jobs to help the scheduler do job scheduling, which can reduce job completion time and cross-datacenter bandwidth. The SAVI testbed is a distributed application platform that allows for data processed across geo-distributed datacenters and thus provides an ideal environment for evaluating the performance of our framework.

▪ *Optimizing Coflow Completion Times with Utility Max-Min Fairness*

Li Chen, Wei Cui, Bo Li, **Baochun Li**, (University of Toronto)

In data parallel frameworks such as MapReduce and Spark, a coflow (<https://amplab.cs.berkeley.edu/tag/coflow/>) represents a set of network flows used to transfer intermediate data between successive computation stages for a job. The completion time of a job is then determined by the collective behavior of such a coflow, rather than any individual flow within, and influenced by the amount of network bandwidth allocated to it. Different jobs in a shared cluster have different degrees of sensitivity to their completion times, modeled by their respective utility functions.

In this paper, we focus on the design and implementation of a new utility optimal scheduler across competing coflows, to provide differential treatment to coflows with different degrees of sensitivity, yet still satisfying max-min fairness across these coflows. Though this objective can be formulated as a lexicographical maximization problem, it is challenging to solve in practice due to its inherent multi-objective and discrete nature. To address this challenge, we first divide the problem into iterative steps of single-objective sub-problems; and in each of these steps, we then perform a series

of transformations to obtain an equivalent linear programming (LP) problem, which can be efficiently solved in practice. To demonstrate that our solutions are practically feasible, we have implemented it as a real-world coflow scheduler on the SAVI platform, based on the Varys open-source framework (<http://varys.net/>), to evaluate its effectiveness.

▪ ***Virtual Network Function Description, Publication and Discovery***

Sandhya Shanmugasundaram, Sami Yangui, **Roch Glitho**, (Université Concordia University)

This poster proposes a novel description model for Virtual Network Functions (VNFs). It also proposes an architecture for VNFs publication and discovery. The current discovery mechanism we have prototyped is based on syntactic matching. It is a two level process. We first try to discover from local repositories and if the discovery is not successful, we then try to discover from external repositories. An illustrative use case is provided using VNFs in content delivery networks (CDNs) settings.

▪ ***A Platform as a Service for Content Delivery Networks***

Aida Rangy, Sami Yangui, **Roch Glitho**, (Université Concordia University)

Platform as-a-Service (PaaS) could be an effective mean for provisioning content delivery network (CDN) applications such as social TV in cloud based CDN settings. Unfortunately, the studies we have conducted on existing PaaS solutions show that none of them is able to easily provision CDN applications. This poster proposes architecture for effectively provisioning CDN applications in cloud settings. Social TV is used as an illustrative use case.

▪ ***Real-time Traffic-based Routing, based on Open Data and Open-Source Software***

Diego Serrano, Teresa Baldassarre, **Eleni Stroulia**, (University of Alberta)

The emergence of cloud computing and the Internet of Things (IoT) have given rise to a wealth of new opportunities for integrating heterogeneous systems and collecting massive data sets whose analysis may lead to new information, insight and knowledge. Building a scalable architecture for urban IoT environments is a complex task, primarily because of the massive amounts of data generated by sensor devices and the variety of data sources. And yet, it is a compelling application area given the number of potential municipal services that can be improved using these technologies.

In this poster, we describe our study of how cloud-computing and big-data management technologies can assist decision making for transportation systems in smart cities. More specifically, this paper presents and discusses a proof-of-concept prototype, based on open-source technologies and publicly available data for the city of Edmonton.

▪ ***Kaleidoscope: A Cloud-Based Platform for Real-Time Video-based Interaction***

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 advancement of wireless networks and mobile technologies.

In this poster, we present, Kaleidoscope, a cloud platform for multimedia streaming on mobile devices, enhanced with textual and touch-display interactions for a rich user experience. We evaluate the Kaleidoscope system on two different clouds at different locations, testing the streaming quality, and CPU and memory usage, using different numbers of clients.

▪ ***PrIMoR-sky: A Runtime Models Processing Infrastructure for Cyber Physical Applications in the Cloud***

Lorena Castaneda, Norha Villegas, **Hausi Müller**, (University of Victoria)

PrIMoR-sky is a smart cloud infrastructure for runtime models management. PrIMoR-sky services include add, delete and update runtime models at execution time. With PrIMoR-sky, cyber physical applications can outsource runtime models adaptation requirements.

▪ ***Green Data Centers: Smart Applications and Energy Metrics***

Andreas Bergen, **Hausi Müller**, **Sudhaker Ganti**, (University of Victoria)

Understanding application specific energy consumption of software components within an Edge/Core architecture allows us to optimize software components dynamically in place, as well as to migrate components within and across clouds. Dynamic energy consumption and optimization of software applications in SAVI has the goal of lowering energy consumption and improving software performance.

▪ ***CAVA: Context Aware Video Analytics on the SAVI Network***

Nina Taherimaksousi and **Hausi Müller**, (University of Victoria)

In this poster we present a context aware video analytics system for continuous collection of first person video from smart mobile devices. It achieves scalability with decentralized cloud on the SAVI network. Videos are processed based on contextual information such as location, time and video context, and run on virtual machines on the SAVI edges. Smart mobile device users can accomplish searches on the context aware processed video database. This poster provides an understanding of the video parameter impact, such as size, speed and resolution, on the scalability of our approach. We also present selected challenges for video collecting, processing, labeling and searching.

➤ **POSTER & DEMONSTATION:**

▪ ***Online Group Chat with Bellini on the SAVI Testbed***

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

Online group chat is one of the SAVI smart applications with demands on low-latency data delivery. A message in a group chat is sent out at a random time, typically with a random size. Also, the messages may include multimedia contents such as video clips and voice messages, which are required to be delivered with a reasonable order. In this demo, we will introduce how Bellini, our effective, transparent and flexible data delivery service can support online group chat applications on the SAVI testbed.

Each web-based group chat client communicates directly with an assigned Bellini server located in the nearest SAVI region. For a given group, the messages will be automatically broadcast to all peers in different regions, where multicast data delivery is enabled. Furthermore, even though the

delivery of chat messages shares the network links with other traffic-intensive applications, they are prioritized such that the data delivery latencies are reasonably low.

▪ *Cross-Region Data Delivery Service on the SAVI Testbed*

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

The SAVI testbed provides numerous and considerable computation and storage resources across geographically distributed regions. Although different regions are interconnected by dedicated links or tunnels, coordinating these geo-distributed resources with cost-effectiveness is a critical problem. In this demo, we present Bellini, an effective, transparent, and flexible data delivery service, which aims at providing applications with a high-level abstraction of cross-region transport.

Bellini employs several worker nodes in each SAVI region, establishing parallel TCP connections and building a cross-region network overlay. Applications are allowed to read from and write to any worker node to access the data delivery service, caring nothing about the transport detail. Under the “hood”, Bellini manages the overlay as a software-defined network, where all traffic among worker nodes is administrated by a centralized controller. This design enables highly configurable policies in data delivery, which can be customized for applications' demands.

In this demo, we will use online instrumenting tools to visualize the operation of Bellini in real time. Specifically, when supporting multiple traffic-intensive applications across SAVI regions, we will show that Bellini can allow them to share the infrastructure effectively. These applications include video broadcast, online group chat with video messages, and data analytic jobs.

▪ *An NFV Based Architecture for On-the-fly Provisioning of Value Added Video Services in Content Delivery Networks*

Elaheh Narjes Taghigh Jahromi, Sandhya Shanmugasundaram, Aida Rangy, Sami Yangui, Mohammad Ali Salahuddin, **Roch Glitho**, (Université Concordia University)

A Content Delivery Network (CDN) is an effective solution for provisioning content services, such as video services, to a large number of geographically distributed users. However, provisioning on the fly value added video services remains a challenge that Network Function Virtualization (NFV) and Software-Defined Networking (SDN) could aid in tackling.

This poster (which comes with a demo) proposes a business model and an architecture for provisioning on-the-fly value added video services in CDN. The service functions that make these value-added services are packaged and deployed as virtual network functions (VNFs) and are chained using application level switches (instead of the traditional IP level switches). This application level chaining is enabled by an extended SDN controller. A software prototype is built using HTTP proxies and concretely demonstrated by a use case on video overlay value added service.

▪ *A Cloud Platform-as-a-Service for Multimedia Conferencing Service Provisioning*

Ahmad Ferdous Bin Alam, Abbas Soltanian, Sami Yangui, Mohammad Ali Salahuddin, Halima Elbiaze, **Roch Glitho**, (Université Concordia University)

Multimedia conferencing is the real-time exchange of multimedia content between multiple parties. It is the basis of a wide range of applications (e.g., multimedia multiplayer game). Cloud-based provisioning of the conferencing services on which these applications rely will bring benefits, such as easy service provisioning and elastic scalability. However, it remains a big challenge.

This poster (which comes with a demo) proposes a PaaS for conferencing service provisioning. The proposed PaaS is based on a business model from the state of the art. It relies on conferencing multiple IaaS that, instead of VMs, offer conferencing substrates (e.g., dial-in signaling, video mixer and audio mixer). The PaaS enables composition of new conferences from substrates on the fly. This has been prototyped and in order to evaluate it, a conferencing IaaS is also implemented. Performance measurements are also made.

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

Using the SAVI infrastructure we are able to dynamically customize 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.

THEME 2 - RESOURCE CONTROL & MANAGEMENT

➤ **POSTERS:**

▪ ***SpearDB: Database Management at the Network Edge***
Cătălin Avram, Bernard Wong, Ken Salem, (University of Waterloo)

A challenge faced by modern cloud applications is the request latencies that are introduced by the geographic separation between clients and remote application servers. One way to address this challenge is to push all or part of the application (along with the data it requires) to the edge of the network - closer to application clients.

SpearDB is a replicated distributed database system designed to operate in the setting of a star network topology, with a core site and a large number of edge sites that are close to clients. Clients access the nearest edge, which holds replicas of locally relevant portions of the database. SpearDB's edge sites coordinate through the core to provide a global transactional consistency guarantee (parallel snapshot isolation), while handling as much work locally as possible.

This poster provides a general overview of SpearDB, focusing on the transactional aspects of the system; along with experimental results showing SpearDB's effectiveness at reducing request latencies for applications whose access patterns are geographically localizable.

▪ ***Modeling Performance Impact of Middleware in Cloud Systems***
Adnan Faisal, Dorina Petriu, Murray Woodside, (Carleton University)

A cloud application may use different kinds of middleware, or different middleware configurations. Choices may have to be made between middleware alternatives that may have a substantial impact on the performance (such as response times or throughput capacity) of the application. This research explores the various middleware decisions by modeling their impact on the application performance, in a Layered Queueing (LQ) modeling framework which reflects the cloud environment

as well as the application organization and workload. It can be used with LQ models that have been estimated using semi-automated modeling tools on a prototype application.

The modeling framework adds and configures the middleware in the model, in a way that allows many deployment and configuration alternatives to be compared. The middleware options are structured with a meta-model incorporating the common operations of middleware products, such as flow encoding and target discovery, and a feature model. A systematic process for defining options and features and building the performance model has been created, with a software tool, and demonstrated on cases built with different kinds of middleware such as Distributed Object Based, Container Based and Message Based middleware.

Cloud computing gives the opportunity to deploy many virtual computers with the same configuration. This feature makes it easy to model and calibrate a middleware once and then reuse it in various applications. This aspect of reusability is one of the major goals of the proposed framework.

▪ *LEGIS on Scale with Docker Containers*

Marios Fokaefs, Cornel Barna, Rodrigo Veleda, **Marin Litoiu**, (York University)

The Locality Enhanced Geographic Information System (LEGIS)* is a distributed navigation service which takes advantage of multi-tier clouds like SAVI to distribute traffic and data analytics. In this work, we focus on the management of a LEGIS deployment. We outline the differences and advantages of various topologies based on virtual machines, docker containers or combinations. We detail the monitoring challenges around all resources and how they have to be analyzed in order to make scaling decisions during workload fluctuations. Finally, we discuss what actions are possible and appropriate to scale such multi-layer, multi-tier and distributed topologies. We present an exploratory process to identify the pitfalls, the challenges and best practices in managing such applications. The system is evaluated on the SAVI infrastructure using real data from the CVST project. *LEGIS won the first place in the SAVI Design Challenge Camp in 2015 among 6 contenders.

▪ *On Efficiency and Scalability of Software-Defined Infrastructure for Adaptive Applications*

Nasim Beigi-Mohammadi, Hamzeh Khazaei, Mark Shtern, Cornel Barna, **Marin Litoiu**, (York University)

This work proposes and evaluates a novel analytical performance model to study the efficiency and scalability of software-defined infrastructure (SDI) to host adaptive applications. The SDI allows applications to communicate their adaptation requirements at run-time. Adaptation scenarios require computing and networking resources to be provided to applications in a timely manner to facilitate seamless service delivery. Our analytical model yields the response time of realizing adaptations on the SDI and reveals the scalability limitations. We conduct extensive testbed experiments on SAVI to verify the accuracy and fidelity of the model. Cloud service providers can leverage the proposed model to perform capacity planning and bottleneck analysis when they accommodate adaptive applications.

▪ *An Adaptive Pricing Model for Scalable Cloud Applications*

Marios Fokaefs, Cornel Barna, **Marin Litoiu**, (York University)

Modelling revenue generation and deciding on pricing policies can prove rather challenging for web software systems. This is primarily because such environments are quite volatile, on one hand in terms of the incoming traffic and on the other hand because the cost for resources is equally volatile due to demand fluctuations.

In this work, we present a demand model that combines technical aspects of the application, e.g. the system's response time, the number and type of the used resources, and the arrival rate, along with economic indexes, such as price elasticity, user satisfactions, and resource costs and so on. This model will assist software providers to optimally price their system, by employing an adaptive and scalable advertisement strategy, in the face of fluctuating traffic and scalable cloud topology. The model is to be integrated as a decision mechanism in a broader autonomic management system for cloud applications and evaluated as a service on the SAVI cloud.

▪ *Provisioning Performance of Cloud Microservices*

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

Microservice architecture has started a new trend for application development or deployment in cloud due to its flexibility, scalability, manageability and performance. These qualities are made possible by the recent advances in Linux containers and by the need to increase resource utilization in cloud. Various microservice platforms have emerged to facilitate the whole cycle of software engineering for cloud applications from development, test, and deployment to maintenance. Microservice platforms rely on macroservice cloud infrastructure in order to acquire the necessary resources.

In this paper, we propose a performance analytical model and validate it by a simulation model to study the provisioning performance of microservice platforms. In order to actualize these models, we design and develop a microservice platform on SAVI cloud to first validate the analytical/simulation models and second to provide exogenous parameters for the performance models. Our results show that container placement strategy, virtual machine de/provisioning, VM consolidation and container live migration have huge impact on provisioning performance of microservice platforms.

▪ *Overload Management for Busy Multi-tier Web Applications using Software Defined Networking*

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

In this research we propose and implement an SDN-based solution to manage overloaded web applications in cloud centers. The idea is to use flow control policies to reduce the loads on the busy application nodes in order to improve the overall response time. To this end, we develop an SDN controller application that controls and modifies the QoS of application flows on the fly. Assigning on-demand bandwidth to the flows is among the implemented control actions. Optimization techniques are used to derive the best place to impose the bandwidth variation on different flows from multiple applications that share the same infrastructure. We use the SAVI testbed to implement the solution.

➤ **POSTER & DEMONSTRATION:**

▪ *Bitnobi - A Big Data Platform for Distributed Clouds*

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

Bitnobi is a Big Data platform which simplifies the development and execution of complex big queries into drag and drop activities that are suitable for both non-technical and technical users. It also solves the main concern on security of data sharing and data segmentation. Our user access control framework allows easy yet secured approach using user attributes as policy to determine user's access permissions. Bitnobi is deployed using SAVI resources and we demonstrate the platform's capabilities using data from the "Connected Vehicles Smart Transportation" (CVST) project.

▪ *Hogna as a DevOps Framework for Containerized Cloud Applications*

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

Recent advancements in software delivery and virtualization of computing and networking resources have reduced software into even smaller nuggets of functionality (microservices) and have raised infrastructure to even higher levels of abstractions (containers). Containers package microservices into isolated software environments (libraries, configurations, file systems etc.) that are hosted by an operating system inside a virtual machine. Containers constitute a more lightweight hosting environment that lifts from the software developer the responsibility of managing and maintaining the infrastructure or even the operating system, bringing development closer to the DevOps model.

We extend our autonomic management system, Hogna, which is already implemented to work with SAVI applications, to enable the management of containerized topologies. We outline how Hogna can monitor containers, VMs and the deployed applications and how it can enable the systematic or even automatic adaptation and scaling of these resources. We present this extension of Hogna in a variety of applications deployed in the SAVI cloud.

➤ **DEMONSTRATIONS:**

▪ *Real-time Detection and Identification of Events on Ontario Highways*

Hamzeh Khazaei, Rodrigo Veleda, **Marin Litoiu**, (York University)

In this demo we present the capabilities of Sipresk platform (<https://tinyurl.com/zu5fz94>) in detection and identification of events in Ontario highways in real-time. Sipresk is a big data analytic platform that is able to perform real-time and retrospective analysis on transportation data.

We design couple of signatures that will be applied on streamed sensor traffic data and generate appropriate alarms in case of events in Ontario highways with the maximum delay of one minute. Each signature has been designed specifically for certain type of events (i.e., accidents and long term road blocks). The platform is cluster-based and leverages the cloud to achieve reliability, scalability and adaptivity to cope with the changing environment. The primary source of data for Sipresk platform is Loop detector data that are being sampled every 20 seconds.

THEME 3 - SMART CONVERGED EDGE

➤ POSTERS:

▪ *Benchmarking and Evaluating OpenFlow Controllers*

Mohamad Darianian and **Carey Williamson**, (University of Calgary)

In network traffic engineering, network administrators need automated mechanisms to detect, classify, manage, and dynamically re-route an evolving set of traffic flows across many possible network paths. In an SDN-based network, this network traffic management functionality requires a high performance and responsive OpenFlow controller.

In this research project, we will use an experimental SDN testbed to evaluate two different OpenFlow controllers (ONOS and OpenDaylight) with respect to controller response time, throughput, and switching delay. The experimental results will help determine which controller is best-suited to near-real-time control of traffic flow management.

▪ *Towards a Scalable Monitoring Infrastructure Using SDN*

Carey Williamson and **Majid Ghaderi**, (University of Calgary)

In collaboration with the City of Calgary, we are currently seeking NSERC SPG funding for an SDN-based project on wireless sensor networks for environmental systems monitoring. The primary objective of the project is to design and deploy a scalable monitoring infrastructure for Internet of Things (IoT) applications.

The proposed platform will consist of hardware/software ICT components that are integrated and networked using SDN to provide a flexible and powerful framework for near-real-time monitoring and control. The initial focus will be on networking weather station data sites for water monitoring and smart sewer applications, though the design will also support other possible target domains (e.g., smart data centers, pipeline monitoring and transportation systems). If this project gets funded, we will need HQP with SDN expertise in Fall 2016.

▪ *Evaluating RTE and NC as Network Services*

Feifei Shi and **Carey Williamson**, (University of Calgary)

Recent estimates suggest that over 90% of current Internet traffic is redundant, with repeated transfers of identical or highly similar content between many different network endpoints. Redundant Traffic Elimination (RTE) and Network Compression (NC) are two possible approaches for alleviating this problem. RTE focuses on detecting and removing repeated chunks of data within packet-level transmissions, replacing them instead with compact fixed-size fingerprints. NC uses information-theoretic approaches to reduce the number of bits required to transmit arbitrary packet data.

In this work, we implement and evaluate both RTE and NC as potential SDN network services. Our evaluation results on empirical email traffic datasets indicate that bandwidth savings of 40% - 70% are possible using RTE and NC.

▪ *Distributed Service Function Chaining*

Milad Ghaznavi, Nashid Shahriar, Shahin Kamali, Reaz Ahmed, **Raouf Boutaba**, (University of Waterloo)

Middleboxes have become a vital part of modern networks by providing service functions such as content filtering, load balancing, and traffic optimization. An ordered sequence of middleboxes composing a logical service is called a service chain. Service Function Chaining (SFC) enables us to define these service chains. Recent optimization models of SFCs assume that a middlebox's functionality is provided by software appliances, commonly known as Virtual Network Functions (VNFs), deployed in a single physical machine. This assumption limits SFCs to the throughput of few VNF instances that can be installed in one physical machine. Moreover, typical service providers offer VNFs with heterogeneous throughput and resource configurations. Thus, deploying a service chain with custom throughput can become a tedious process of stitching heterogeneous VNF instances.

This work describes how to overcome these limitations without worrying about underlying VNF configurations and resource constraints. Doing so is achieved through the distributed deployment of multiple VNF instances that provide the functionality of a middlebox and modeling the optimal deployment of a service chain as a Mixed Integer Programming problem (MIP). The proposed model optimizes host and bandwidth resource allocation, and determines the optimal placement of VNF instances, while balancing workload and routing traffic among these VNF instances. We prove that this problem is NP-Hard and propose a heuristic solution called Paliz. Paliz utilizes a tuning parameter to control the trade-off between speed and accuracy of the solution. Finally, our solution is evaluated using simulations in data-center networks.

▪ *Fabric Embedding with Link Failure Protection Guarantees*

Aimal Khan, Nashid Shahriar, Milad Ghaznavi, Shihabur Rahman Chowdhury, Reaz Ahmed, **Raouf Boutaba**, (University of Waterloo)

Evolution in virtualization technologies and diversity of customer demands are changing the way Virtual Networks are provisioned and used. It is now possible for an Infrastructure Provider (InP) to control and manage its networks more flexibly, at finer granularities. This gives rise to use-cases where customers demand for easier management of their virtual resources and flexible service models. The InP can have tenants that are not interested in managing their virtual networks (VNs), and require a managed 'Network Fabric' as a service. In such a service model, the tenant is only required to provide traffic demands between selected virtual resources, and the onus of designing, provisioning (mapping on the physical infrastructure), providing QoS guarantees (e.g. service availability) and managing the lifecycle of the VN is on the InP. In this context, the objective of an InP is to provide efficient fabric embedding while minimizing the total provisioning costs.

In this work, we look at the problem of providing Virtual Network Fabric services with service availability guarantees from the InP perspective. More specifically, we look at the problem of Embedding Virtual Network Fabrics that are resilient to single physical link failures. We provide a mathematical model that considers the problems of Fabric Embedding and Fabric Survivability in a joint manner. The model provides optimal placements for VN resources (nodes and links) while minimizing the cost of providing survivability guarantees in case of a single physical link failure. Moreover, we propose a fast heuristic solution that scales with larger sized networks while keeping within acceptable error margins of the optimal solution.

▪ *Connectivity-aware Virtual Network Embedding*

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

The problem of ensuring virtual network (VN) connectivity in presence of multiple link failures in the substrate network (SN) is not well investigated in Network Virtualization (NV) literature. We name this problem as Connectivity-aware Virtual Network Embedding (CoViNE). Solving CoViNE will enable a VN operator to perform failure recovery without depending on the SN provider, similar to the IP restoration mechanisms in IP-over-WDM networks. There are two steps in solving CoViNE: i) finding the virtual links that should be embedded disjointly, and ii) finding a substrate resource efficient embedding that ensures the virtual link disjointness constraint.

We present two solutions to the CoViNE problem. The first solution uses a heuristic to compute the disjointness constraint, while an optimization model is used for VN embedding. The second solution, in contrast, uses heuristics for both steps, and thus can solve larger instances of the problem. We compare our solutions with a cut set based approach that ensures VN connectivity for a single substrate link failure. Evaluation results show that our heuristics allocate 15% extra resources on average compared to the cut set based optimal solution, and executes two to three orders of magnitude faster on the same problem instances.

▪ *On Orchestrating Virtual Network Functions*

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

Middleboxes or network appliances like firewalls, proxies, and WAN optimizers have become an integral part of today's ISP and enterprise networks. Middlebox functionalities are usually deployed on expensive and proprietary hardware that require trained personnel for deployment and maintenance. Middleboxes contribute significantly to a network's capital and operational costs. In addition, organizations often require their traffic to pass through a specific sequence of middleboxes for compliance with security and performance policies. This makes the middlebox deployment and maintenance tasks even more complicated. Network Function Virtualization (NFV) is an emerging and promising technology that is envisioned to overcome these challenges. It proposes to move packet processing from dedicated hardware middleboxes to software running on commodity servers.

In NFV terminology, software middleboxes are referred to as Virtual Network Functions (VNFs). It is a challenging problem to determine the required number and placement of VNFs that optimize network operational costs and utilization, without violating service level agreements. We call this the VNF Orchestration Problem (VNF-OP) and provide an Integer Linear Programming (ILP) formulation with implementation in CPLEX. We also provide a dynamic programming based heuristic to solve larger instances of VNF-OP. Trace driven simulations on real-world network topologies demonstrate that the heuristic can provide solutions that are within 1.3 times of the optimal solution. Our experiments suggest that a VNF based approach can provide more than 4x reduction in the operational cost of a network.

▪ *Revolutionizing Home Firewalls with FleXam*

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

The Firewall is an important part of network security and protection. Current firewalls are complex systems that are neither easy to manage nor cheap, making it difficult for home network owners to operate and maintain them, which results in less secure networks.

In this paper, we review different enterprise firewall designs and discuss why they are not suitable for home networks. Then we propose a platform to implement software defined firewalls that enables home network owners to delegate part of network monitoring to a third party service provider. The core of our system is a flexible sampling extension for OpenFlow that can intelligently send a subset of packets that are required for decision making to the remote controller. This gives the controller an in-depth knowledge of network traffic (even full packet contents) with a manageable overhead.

▪ *Dependence of TCP Performance on Congestion Control Parameters*

Shiva Ketabi, Monia Ghobadi, Soheil Hassas-Yeganeh, Yashar Ganjali, (University of Toronto)

Transmission Control Protocol (TCP) has been dominantly used in computer networks during the last few decades, despite its known shortcomings. Researchers have tried to improve TCP's performance and to adjust its parameters. Their solutions are mostly domain specific and cannot be generally applied. Also, they are static as they cannot be updated during the run-time.

OpenTCP is a dynamic framework that can monitor the network and accordingly modify TCP's behaviour. In this framework, switches collect the data from the network, a controller receives the data and generates congestion control decisions, and end-host agents change the end-host's TCP based on the controller's commands. We investigate the effects and challenges of tweaking the aggressiveness of TCP's Additive Increase Multiplicative Decrease (AIMD) scheme using a static framework. We repeat the experiments for the dynamic adaptation framework. Also we tune the initial congestion window (icwnd) and compare the results of a static framework with OpenTCP. The experiments reveal two major points: first, it is worth making the effort to adjust TCP parameters since it can significantly improve TCP performance measures. Second, a dynamic framework can help to modify TCP parameters in a controlled manner and can guarantee that the undesirable side effects of tweaking TCP are eliminated.

➤ **DEMONSTRATIONS:**

▪ *NF.IO: Network Function Discovery and Orchestration*

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

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. We propose a NF discovery and orchestration system called NF.IO that adopts these principles. It uses the standardized and well-known POSIX compliant file system interface for NF orchestration. We have developed a prototype, and in this demonstration we will showcase the capabilities of NF.IO through the following use cases: (i) configuration, (ii) deployment, and (iii) monitoring.

➤ POSTERS:

▪ *A New MAC Protocol Design to Support Resource Management in Virtual Wireless Networks*

Daniel Tweed, Alfred Kenny, Nhat-Quang Dao, Quang-Dung Ho, **Tho Le-Ngoc**, (McGill University)

Wireless networks have developed to suit many diverse use-cases. This has resulted in myriad, often divergent, approaches to medium access, limiting their widespread applicability. Distributed and stochastic protocols are good at handling diverse network architectures and traffic profiles, but are not able to guarantee QoS or slice isolation in virtualized networks. Centralized and deterministic approaches can provide these guarantees, but are inefficient in sparse non-saturated networks. Further, they require a level of synchronization that can be both difficult to achieve and to maintain.

To address these challenges, we propose DFRA, an asynchronous, centrally scheduled MAC protocol built upon 802.11 Wi-Fi as an enhancement to CSMA/CA. Station schedules are embedded into each beacon frame, with the inter-beacon interval sub-divided into resource blocks which can be allocated to specific users or contended for amongst a group of users. DFRA supports scheduling algorithms in a plug and play manner and can dynamically adapt to changing traffic requirements to support diverse heterogeneous networks. Thus, DFRA can benefit from both strict scheduling and random access, as needed, and can provide guaranteed QoS and slice isolation in virtualized networks.

▪ *Inter-Access Point Coordination for Efficient Network Resource Management*

Nhat-Quang Dao, Daniel Tweed, Alfred Kenny, Quang-Dung Ho, **Tho Le-Ngoc**, (McGill University)

In dense wireless access network deployments, uncoordinated Access Points (APs) can introduce heavy channel contention and interference that may hinder the capability to provide priority medium access of their respective users. Coordination between APs is one of the promising approaches to efficiently allocate resources among APs and facilitate priority scheduling of stations associated with different APs.

This work describes a network structure with associated communication protocols that support time-slotted scheduling for priority medium access of multiple APs and their associated stations through a Central Controller (CC). Information is sent to CC from APs to report their number of associated stations, number of detected neighboring APs, and any priority access requests their stations have made. With such global information available to the CC, the framework enables cooperative and dynamic resource allocation (such as time slot scheduling, access priority assignment, and frequency selection) for efficient resource utilization and differentiated QoS provisioning. The communication requirements and CC operations necessary to coordinate between APs and CC are investigated and the performance is evaluated with an illustrative MAC protocol, namely Dedicated, Free-assigned, and Random Access (DFRA).

▪ *SiP Modulators for Radio over Fiber*

Jiachaun Lin, Bahareh Sherafati, Hassan Sepehrian, Wei Shi, **Leslie A. Rusch**, (Université Laval)

We present a modulation power efficient silicon Mach-Zehnder modulator for radio signal delivery in dense small cell access point. The silicon modulator is designed with forward-bias PN junction. Both of its DC and AC characteristics are studied. A very low figure of merit of $V_{\pi}L$ is achieved compared to a commercialized LiNbO₃ MZM. With the designed low power consumption silicon MZM, we successfully delivered 4Gb/s OFDM signal through 20km fiber link in an intermediate frequency over fiber (IFoF) system.

▪ ***Constellation-Induced SNR Gain in Short-Reach Optical OFDM***

Mingyang Lyu, Siamak Amiralizadeh, **Leslie A. Rusch**, (Université Laval)

We investigate constellation-induced SNR gain in short-reach OFDM with direct detection for various distances and laser linewidths. Performance improvement by using alternate constellations is more significant for lower quality lasers, especially for 8-QAM. As transmission distance increases, higher SNR gain is obtained (1.92dB for 100km).

▪ ***Closing the Loop: Data and RoF Symbiosis in Uplink and Downlink***

Zhihui Cao, Jiachuan Lin, **Leslie 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). Virtualized WiFi access point integrated with our transparent radio over fiber system has been experimentally demonstrated. The throughput up to 54Mbps (802.11g) at the RAU could be achieved with full duplex error free OOK transmission over up to 20 km single mode fiber. While negligible penalty (less than 1Mbps drop) introduced by the radio over fiber link to the access point throughput has also been demonstrated.

➤ **POSTER & DEMONSTRATION:**

▪ ***Bringing Control of Radio Resources to Virtualized Networks: The Implementation of a Programmable Medium Access Protocol***

Alfred Kenny, Nhat-Quang Dao, Daniel Tweed, Quang-Dung Ho, **Tho Le-Ngoc**, (McGill University)

Effective virtualization in wireless access networks requires that tenants of these networks have abstracted control of virtual network configurations and resources and that any changes in functionality in one virtual network do not affect the operations of other networks running on the same hardware. This work presents the architecture and implementation of a programmable medium access protocol as a component of an overall virtualization platform. This platform allows network tenants to create/delete virtual wireless networks, and to configure the wireless resources in a cluster of access nodes of the same virtual network via inter-node communications.

The medium access protocol makes use of slotted, carrier sense-based scheduling to produce a time-based resource unit that can be dynamically programmed for demand-assignment, free-assignment, or random-access (DFRA) among specific clients of virtual networks. This resource programmability can improve the efficiency of resource usage, and support the reconfiguration and coordination among virtual networks. Through the implementation of DFRA in a real world environment, the

challenges and opportunities in developing medium access virtualization are highlighted and the efficacy of DFRA for wireless access networks is presented.

THEME 5 - APPLICATION PLATFORM TESTBED

➤ POSTER:

▪ *Scalable Reconfigurable FPGA Clusters in the Cloud*

Naif Tarafdar, Erik Fukuda, Jiayuan Chen, **Paul Chow**, (University of Toronto)

We present infrastructure to generate scalable FPGA clusters within the SAVI cloud infrastructure. This provides the user a unified view of the fabric. This infrastructure allows the user to create connected FPGA circuits across multiple FPGAs where we abstract away the inter-FPGA network connections, scheduling for master-worker models, and other scheduling algorithms between FPGA cores. This infrastructure can be used to target multiple families of applications ranging from master-worker model applications such as web search, and streaming applications such as Network Function Virtualization and media streaming/processing applications.

▪ *Organizing FPGAs in the Cloud for Distributed Applications*

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

FPGAs are making their way to data centers and more application developers are starting to get access to FPGAs. However, methodology for organizing distributed FPGAs in data centers is not established yet. In our work, we try to organize distributed FPGAs as well as processors with the help of existing distributed software applications. These applications include YARN (datacenter resource management software), ZooKeeper (a distributed task management software), and Drill, (a distributed SQL processing software). We use FPGAs for improving the performance of filtering operations of Drill, leaving other operations to distributed processors.

▪ *Autonomous Power Consumption Scheduling for Smart Homes with Local Power Generation and Storage Facilities*

Mohammad Hossein Yaghmaee Moghaddam, Morteza Moghaddassian, **Alberto Leon-Garcia**, (University of Toronto)

Demand Side Management (DSM) is an important application of the future Smart Grid (SG). Generally, DSM programs vary in nature but the aim of such programs is to enable customers to participate in the operation of the electric grid by reducing or shifting their electricity usage during peak hours to avoid power shortage and to decrease the grid operational costs. Direct Load Control (DLC), Time-of-Use Pricing and Autonomous Demand Side Management are examples of such programs. In autonomous models, usually the demand side management program schedules the hourly demand levels of the grid customers to reveal the pressure on power generators, especially during peak hours. However, for the purpose of scheduling, the demand side management program needs electric devices to be “shiftable” (schedulable).

In our model, we also define an autonomous scheduling model to minimize the operational cost of the grid in which the customers are not only equipped with “shiftable” electric devices but are

also utilized with local power generation and storage facilities, that is most likely to happen in future smart homes.

▪ *Graph-based diagnosis in Software-Defined Infrastructure*

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

Performing system diagnosis is a critical task in modern datacenters. Investigating individual resource behavior may not be efficient in detecting abnormal behavior in large and complex datacenters.

In this paper, we propose a scalable graph based diagnosis framework to detect system anomalies in Software-Defined Infrastructure running in SAVI testbed. We have leveraged Graph Mining and Machine Learning techniques in our approach in order to detect different kinds of anomalies. We have tested our framework on several use cases: Webserver-Database workload pattern, bandwidth throttling between a pair of VMs, denial-of-service (DoS) attack on a webserver and Spark Job failure. Our framework was able to detect the aforementioned anomalies accurately.

▪ *Enhanced Real Time Content Delivery using vCPE and NFV Service Chaining*

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

Real-time content delivery (RTCD) systems have become a prominent aspect of telecommunications as evidence by popularity of news-casting, real-time event subscription / publication and live media streaming. Unlike conventional content delivery systems, RTCDs need to deliver processed information to users in real time. This may require the network to handle some of the processing closer to the users to efficiently use the bandwidth consumed by the applications. The combination of Network Function Virtualization (NFV) and service chaining is a promising solution to address this challenge.

Our work applies a service chaining algorithm to place NFV modules of an RTCD application in a Software Defined Infrastructure (SDI), where virtualized Customer Premise Edges (vCPEs), possessing scarce resources, are employed. We suggest containers to efficiently pack VNFs into vCPEs. Our objective is to maximize the total number of chains that can be serviced in the RTCD application. To optimally chain the NFV modules, a heuristic algorithm is proposed and evaluated. Using simulations, we show that our algorithm with the help of vCPEs can support higher number of users while providing high-level service quality.

▪ *OpenFlow Conflict Detection and Authorization in Multi-Tenant Clouds*

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

As a future applications platform, one of the goals of the SAVI testbed is to be flexible enough such that it enables researchers to create applications running novel Future Internet protocols. Thus, the SAVI testbed employs the use of Software-Defined Networking, specifically, OpenFlow, which permits the fine-grained network control required by researchers. However, granting researchers the ability to programmatically control their own flows, at granularities as low as Layer 2, exposes the network to misconfiguration errors that can cause failures in end-to-end reachability, access control, and other security mechanisms.

The current OpenFlow protocol does not have any mechanism for detecting or resolving flow conflicts for maintaining the configuration integrity of the network. When multiple OpenFlow table

entries are matchable to a certain flow, the end behaviour is unpredictable and vendor-specific. Therefore, an efficient flow conflict detection technique is needed in the control plane to aid in the authorization of new flow requests. In this work, we present an authorization and authentication mechanism for safeguarding the network from malicious flow requests, ensuring network configuration integrity, and enforcing tenant isolation in multi-tenant cloud environments.

▪ *Cloud Resource Monitoring for Facilitating Administration*

Sumit Kadyan and **Sudhakar Ganti**, (University of Victoria)

This project presents the design and implementation of a web dashboard used for monitoring cloud resources. Cloud computing provides sharing and managing data and performing computations on a shared resource via the Internet. Here the cloud that is being monitored by the dashboard is the SAVI testbed. The dashboard provides a bird's eye view for the monitoring of SAVI test bed resources. It also allows the user to monitor different resources: parameters such as cpu utilization, memory usage, cpu time, virtual cpu and instances. It also allows users to browse through information about different resources and then make a decision on resource migration.

➤ **POSTER & DEMONSTRATION:**

▪ *End-to-End Peering across SAVI and GENI Testbeds using SDI-based Exchange Points*

Saeed Arezoumand, Hadi Bannazadeh, **Alberto Leon-Garcia**, (University of Toronto)

With the growing pervasiveness of virtualization technologies, carrier networks are shifting from simple packet delivery platforms to multi-tenant integrated clouds offering fine-grained service orchestration mechanisms. This will allow third-parties (i.e. tenants) to leverage service provider infrastructures to deliver global-scale services. As a result, end-to-end orchestration will arise (as a challenge) when orchestration has to be done among multiple cloud regions that possibly have different authorization domains and control logics.

In this paper we present an inter-domain orchestration model that leverages SDI-based exchange points – SDIXP – to fill the control logic gaps between autonomous service providers. Our model includes a novel design for exchange point as an autonomous SDI node, offering fine-grained integrated capabilities for Software Defined Networking and Network Function Virtualization. We have presented a design for layer 2 peering between the GENI and SAVI testbeds as a use-case for our end-to-end orchestration model.

▪ *NFV Service Chaining in SAVI Testbed*

Saeed Arezoumand, Thomas Lin, Hadi Bannazadeh, **Alberto Leon-Garcia**, (University of Toronto)

Service Function Chaining is the key requirement of Network function virtualization. Current proposals for service function chaining rely on transport encapsulation (e.g. GRE, VXLAN, MPLS, etc.) which enforces extra overheads, or using Network Service Headers (NSH) which is not fully adapted to current infrastructures.

In this work, we present a Service Chaining Platform for the SAVI testbed. This platform leverages the native forwarding capabilities of the SAVI SDI architecture to realize service function chaining. The SAVI SDI architecture allows users to include heterogeneous resources such as GPUs and FPGAs in service chains to accelerate packet processing. Using this platform, users can define custom

service chains composed of multiple virtualized network functions and steer a specific flow of traffic through these chains. Traffic flows can be specified based on endpoints and ports. As a use-case, users can steer web traffic destined for a specific web server to pass through a chain of firewalls and DPIs.

▪ *Extending Software-defined Infrastructure to Legacy Systems*

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

Software-defined infrastructure (SDI) offers many advanced capabilities over legacy systems, such as heterogeneous resource management, centralized topology view, and software-defined networking. We propose Vino, a system for extending SDI to public clouds and private resource pools. This system provides a unified view over diverse and distributed substrates. In addition, it allows for advanced orchestration scenarios by leveraging SDI features and knowledge of the substrates' capabilities.

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.

