Vehicular Sensing and Control
Table 1

We will demonstrate Vehicular Sensing and Control (VSC) platform that we designed for researchers and experimenters to support idea/protocol experimentation and emulation in a real world. Specifically, by leveraging the GENI resources,(i.e. GENI WiMAX network, ExoGENI racks, VLAN connection as well as relevant experimenter tools), this platform enables the researchers and experimenters to use OpenXC sensing and camera-based sensing data such as vehicle location, vehicle speed, pedal position, traffic states surrounding the vehicle as input for the experimentation and emulation. It could be very helpful to reason about the protocol design. In this demo, we will highlight the latest mechanisms developed for VSC application-layer emulation and extensive experimental results that we collect using the VSC platform. The audience with the interests of vehicular network should stop by and see this demo given the capability of the VSC platform.

Participants:
* yuehua.research@gmail.com

MobilityFirst
Table 2

Mobilityfirst supports flexible extensibility of data-plane services through a pluggable compute-plane that can optionally and strategically be enabled at chosen points in the network. Service providers can use this to provide value-added services such as caching, security, and offloading options to enhance end-user experience. It can also enable an in-network platform (PaaS-like) for third-party web/consumer application providers to host their servers closer to the their clients. In this demo, we will show how a content delivery service can leverage this compute-layer function to introduce smart in-network modules that can either modify end-to-end streams or respond to requests locally reducing load on origin servers. The demo will utilize host and network resources on GENI (Rack nodes, I2 supported multi-site VLAN using AL2S, etc) as well as wireless edge deployments over WiMAX and !WiFi.
Participants:

* Wireless Information Network Laboratory (WINLAB) at Rutgers University
* Franceso Bronzino, bronzino@winlab.rutgers.edu
* Ivan Seskar, seskar@winlab.rutgers.edu
* Kiran Nagaraja, nkiran@winlab.rutgers.edu

GENI Desktop
Table 3

GENI Desktop provides a unified interface and environment for experimenters to create, control, manage, interact with and measure the performance of GENI slices.

This demo will show the newly implemented functions in the GENI Desktop. We enhanced the GENI Desktop to use the Common Federation Service API V2 to provide stitcher support for slices using RSPECs that need layer-2 connections. We also developed a slice membership function that allows users to be added as a member of a slice after the slice/slivers have been created, enabling new users join and access slice resources. We also have new GENI Desktop functionality to provide a routing support function, and a module for handling netflow data. In addition, we will demo the function implemented in the Adopt-A-GENI project for providing support to user-defined SDN routes through the AAG controller. This demo is suitable to GENI experimenters (beginners and experienced) who want to learn how to manage/control their experiments and interact with GENI resources. It may be also interesting to GENI tool developers who want to see how GENI Desktop interacts with other tools such as the GENI portal and Flack.

Participants:
* Jim Griffioen, griff@netlab.uky.edu, Univ. of Kentucky
* Zongming Fei, fei@netlab.uky.edu, Univ. of Kentucky
* Hussamuddin Nasir, nasir@netlab.uky.edu, Univ. of Kentucky

Simulation-as-a-Service App
Table 4

We will demonstrate an extended simulation- as-a-service (SMaaS) App that involves TotalSim using GENI for PaaS experiments, which will enable them to deliver their App (that has data-intensive computation and data movement workflows) in SaaS form to their customers. We will have a new Layer 2 connection setup from demo site to the TotalSim office in Dublin, OH, and we will show collaboration use case on the virtual desktop environment setup in a GENI Rack at the Metro Data Center. We plan to use Ohio Supercomputer Center resources to integrate some computation element in the demo also and use supporting Apps for analysis/visualization of the results. Gigabit App developers and cloud infrastructure engineers will particularly find our demo interesting.

Participants:
* Prasad Calyam, calyamp@missouri.edu, Univ. of Missouri
* Ray Leto, rleto@totalsim.us, TotalSim

WiMAX Connectivity and Monitoring
Table 5

This demo is an ongoing project to GENI enable an ecological science community utilizing the GENI WiMAX facility. Recent additions consist of IP-Passthrough enabled client devices with an integrated Openflow switch and connections to the GENI rack at the University of Wisconsin-Madison.

Participants:

* Wisconsin Wireless and NetworkinG Systems (WiNGS) Laboratory
* Derek Meyer, dmeyer@cs.wisc.edu
* Suman Banerjee, suman@cs.wisc.edu
* Thomas Steele, Kemp Natural Resources Station, Univ. of Wisconsin

Network Prototype Simulator
Table 6

Network Prototyping Simulator is a simulation system that expands Mininet network emulator to computer cluster. That allow us to reproduce the network with such an amount of nodes that hardly was possible before. The maximum size of network topology in NPS depends on number of cluster nodes with Mininet instances. One cluster node can emulate more than thousand hosts, and an modern server could execute at least 15 cluster nodes packed in virtual machines. As the result, we get about 15 thousands hosts per server. The scalability of NPS makes it possible to emulate really big networks.

By the architecture, NPS saves features of Mininet, so it does not become a clear simulation system, it remains a network prototyping system. Means one could trust the results of such simulation and there is no need to prove correctness and adequacy of the model built.

Participants:
* Vitaly Antonenko, VAntonenko@arccn.ru, Applied Research Center for Computer Networks

ARCCN VERMONT
Table 6

One of the aims of network engineering is to configure forwarding rules of the switches as to guarantee network compliance with the Packet Forwarding Policies (PFP). VERMONT is a software toolset that provides some automation to the solution of this task by checking consistency of OpenFlow rule tables with formally specified invariants of PFP.

VERMONT can be installed in line with the control plane to observe state changes of a network by intercepting messages sent by switches to the controller and commands sent by the controller to switches. It builds an adequate formal model of a whole network and checks every event, such as installation, deletion, or modification of rules, port and switch up and down events, against the requirements of PFP. Before a network update command is sent to a switch VERMONT checks in advance the result of its execution. If a new state of network satisfies all requirements of PFP then the command is delivered to the corresponding switch. Upon detecting a violation of PFP VERMONT blocks the update command, alerts a network administrator, and gives some additional information to discover a possible source of an error.

VERMONT has a wide area of applications. It can be attached to a SDN controller (or a bunch of controllers) to enforce basic safety properties (the absence of loops, blackholes, etc) of the subordinate flow-tables in data plane. VERMONT may be also used as a fully automatic safeguard (and, in a sense, a debugger) for any software application implementing certain PFP on a SDN controller. This properties make VERMONT a desired tool for network administrators and control application developers.

Participants:
 * Vitaly Antonenko, VAntonenko@arccn.ru, Applied Research Center for Computer Networks

Chandelle
Table 7

In this demo we present Chandelle system that shows the benefits of having WLAN networks on tops of SDN infrastructure: faster and seamless migration procedure and cost reduction of wireless access point. This demo also discusses the principles of interaction of the wireless controller and the SDN controller and methods to avoid excessive overhead for the user traffic.

Participants:
 * Vitaly Antonenko, VAntonenko@arccn.ru, Applied Research Center for Computer Networks

Demorpheus
Table 7

In this presentation we propose an approach and hybrid shellcode detection method, aimed at early detection and filtering of unknown 0-days exploits at the network level. The proposed approach allows us to summarize capabilities of shellcode detection algorithms developed over the last ten years into an optimal classifier. The proposed approach allows us to reduce total false-positives rate to almost zero, provides full coverage of shellcode classes detected by individual classifiers, and significantly increases total throughput of detectors. Evaluation with shellcode datasets, including Metasploit Framework plain-text, encrypted and obfuscated shellcodes, benign Windows and Linux binaries, random (normal) data and multimedia shows that hybrid data-flow classifier significantly boosts analysis throughput for benign data - up to 45 times faster than linear combination of classifiers, and almost 1.5 times faster for shellcode datasets. We also give a tool demonstration.

Participants:
 * Vitaly Antonenko, VAntonenko@arccn.ru, Applied Research Center for Computer Networks

Enterprise Network Administration
Table 8

We present a new SDN/OpenFlow-based network management system for Enterprise networks called EasyWay. It demonstrates novel approach where network administrators no longer need to manually congure all network devices, they can simple "draw" a path between network elements and the system will automatically program the network elements. The demonstration shows possibility to manage the complex network from nice graphical interface without manual accessing to network elements.

Participants:
 * Vitaly Antonenko, VAntonenko@arccn.ru, Applied Research Center for Computer Networks

NFV
Table 8

The NFV platform is designed to deploy high-performance virtual network services (VNF) on commodity servers and automatic centralized VNFs orchestration.

The main parts of the NFV platform:

 * High performance data path inside commodity servers.
 * Centralized management system orchestrating virtual services (VNFs) across the servers.
 * The pool of VNFs.

Participants:
 * Vitaly Antonenko, VAntonenko@arccn.ru, Applied Research Center for Computer Networks

Video Streaming over ProtoRINA
Table 9

We leverage ProtoRINA, a prototype of our recursive network architecture, to support a video streaming application. In this demonstration, all traffic between a video streaming server and a video player client, both running on legacy hosts, is redirected by RINA proxies to a RINA network, i.e., a Distributed IPC Facility (DIF). We show superior video quality due to dynamic path adaptation inside the DIF in response to node/link failure.

Participants:
 * Ibrahim Matta, matta@bu.edu, Boston University

End-to-End Flow Manipulations
Table 10

Demonstrate how end-to-end flow manipulations can be performed on programmable OpenFlow-based switches. The purpose is to create a framework for network engineers to utilize as a network troubleshooting/debugging tool. Network debugger uses flow separation at aggregation points to increase granularity in monitoring by using end point analysis tools as debugging points. We will create an environment on GENI to show the framework and its proof-of-concept viability as a network troubleshooting tool.

Participants:
* Deniz Gurkan, dgurkan@central.uh.edu, Univ. of Houston
* Ateeth Kumar Thirukkovulur, ateethkumar@gmail.com

Simulation, Detection, and Denial of Ping Attack
Table 10

The demo topology consist of 4 nodes which includes two hosts, Host1 and Host2, one OpenvSwitch (OVS) node and one Floodlight Controller. The hosts are connected to the OVS and OVS is connected to the Floodlight controller. Host1 will perform ping attack to Host2. Host2 will continuously measure the number of pings received from Host1 and detect the ping attack when the frequency and number of pings exceed a threshold. The detection will trigger re-programming of the network through OVS to deny the pinging connection. All ping attack detection should be monitored using LabWiki and the re-programming of the network is through the GIMI framework.

Participants:
* Deniz Gurkan, dgurkan@central.uh.edu, Univ. of Houston

Measurements of Forwarding Performance of the LINCx Switch
Table 11

This project mainly focuses on the performance measurement of the LINCx software switch (https://github.com/FlowForwarding/lincx) as well as investigations for performance improvements in the forwarding systems. LINCx switch has its flow table hardcoded instead of using ETS table as seen in regular LINC switch. Its flow table code is automatically generated based on the flow instruction message from controller, and recompiled as a module to keep the switch operating without interrupting. This demo will measure the forwarding performance of LINCx after recompilation with different sets of flows.

Participants:
* Deniz Gurkan, dgurkan@central.uh.edu, Univ. of Houston
* Kyle Longtran, kyle.longtran@gmail.com

Internet2
Table 12

Participants:
* Eric Boyd, eboyd@internet2.edu, Internet2

SciWiNet
Table 13
The objective of !SciWiNet is to help the academic research community deal with the wireless infrastructure requirements. !SciWiNet is based on an Mobile Virtual Network Operator (MVNO) model. An MVNO is an entity that re-sells cellular services from a wireless network operator. Typically MVNOs cater to focused user communities. For example !VirginMobile and Straight Talk cater to customers that require single devices with low service rates. Other MVNOs support vertical or focused markets such as Kajeet (education). This article provides further background on MVNOs. Currently, !SciWiNet resells Sprint's 3G/4G data services. This includes Sprint's WiMAX and LTE services (including Sprint's current LTE brand referred to as Spark).

Participants:
* Jim Martin, JMARTY@clemson.edu, Clemson Univ.

GENI MOOC
Table 14

We will demo an experiment-based Massive Open Online Course (MOOC) on the subject of computer networks, with lab experiments that run on GENI resources. This course is aimed at beginners who want to learn about how the Internet works, students who want an introduction to some research topics in networking, and instructors who may use these browser-based experiments as in-class demonstrations or homework assignments. The course will be open to the public and is scheduled to run in early fall or late summer; interested participants should stop by to see the demo and learn how to register. We will also show a poster on our WiMAX-based lab experiments, which may be of interest to instructors.

Participants:
* Fraida Fund, ffund01@students.poly.edu, NYU Polytechnic

Teaching & Tmix
Table 15

Stop by this dual demo if you are: (a) an educator who teaches undergraduate networking courses, and are interested in using our education modules -- ready to use in your class, complete with slides! These consist of concept demonstration

modules as well as student assignment modules, and/or (b) an experimenter who would like to generate realistic traffic in your experiments using our Tmix traffic generation system.

Participants:
* Jay Aikat, aikat@cs.unc.edu, Univ. of North Carolina

Application Aware Big Data
Tables 16, 20

Demo description paragraph (three sentence minimum): We demonstrate an application aware big data experiment across multi-domain, optical-wireline-wireless heterogeneous network so that the nature of the Big Data applications should utilize the best resources. A broker-based multi-domain OpenFlow control framework is designed to intelligently control the heterogeneous network. Two 4K cameras and three 4K monitors are deployed to generate the big data applications (over 10Gb/s), which are flexibly routed by controlling the optical switches. Network providers and service providers that support OpenFlow should see this demonstration because it will let them manage multi-domain heterogeneous networks efficiently.

Participants:
* S.J. Ben Yoo, sbyoo@ucdavis.edu, Univ. of California-Davis
* xtfeng@ucdavis.edu, Univ. of California-Davis
* Rob Proietti, rproietti@ucdavis.edu, Univ. of California-Davis
* Lei Liu, leiliu@ucdavis.edu, Univ. of California-Davis

Virtualized Services via SDN
Table 17

This system demonstrates a novel methodology for providing arbitrary virtualized services via Software Defined Networking (SDN). The tabletop demonstration provides virtual channels for services from the service provider all the way through to the user past the user networking hardware. The SDN based system enables dynamic control of priority based packet forwarding and bandwidth allocation.

Participants:
* Dumidu Wijayasekara, dumidu.wijayasekara@gmail.com

Virtual Computer Networking Lab
Table 18

In this demo we will highlight some of the classroom assignment we have created within the scope of our GENI Education Project "Virtual Computer Networking Lab (VCNL)". We will demonstrates two assignments that were used in the "Computer Networks and the Internet" course taught in the ECE department at UMass Amherst

this spring. We will also present two additional assignments that will be made available for classes being taught this coming fall. VCNL makes use of the GIMI tools and !LabWiki is used by the students to carry out the assignments. It is designed such that assignments can be carried out on both the ExoGENI and the InstaGENI testbed.

Participants:
* Mike Zink, zink@cs.umass.edu, University of Massachusetts

GENI-VIRO
Table 19

We will provide demo to illustrate the current progress of our GENI-VIRO implementation progress using the GENI testbed. We will be using the OVS and SDN platform running on virtual machines running GENI nodes in one or multiple sites to run our current GENI-VIRO code. We will show set up a simple topology with a few VIRO switches, and how packets are routed from a virtual host to a VIRO switch to another virtual host attached to another VIRO switch using the VIRO routing protocol based on Virtual Routing Id (implemented by re-using MAC addresses in layer 2).

Participants:
* Zhi-Li Zhang, zhzhang@cs.umn.edu, University of Minnesota

Jacks
Table 21

Jacks is a new tool for creating and viewing GENI slices and slivers. We will demonstrate Jacks in both of these settings, showing how it integrates with other tools in order to provide maximum flexibility. As an editor, Jacks lets you quickly construct an experimental topology that can be used to create a sliver at the GENI Portal, in Aptlab, or using the Omni toolset. We will also show how Jacks can let you see and interact with slivers you have already created.

Participants:
* Rob Ricci, ricci@cs.utah.edu, Univ of Utah

InstaGENI
Table 21

Demonstrate the properties of the InstaGENI rack at GEC-20. Creation of slices on InstaGENI and a demo application will be shown.

Participants:
* Rick !McGeer, rick@mcgeer.com, US Ignite

VNode, FLARE, Congestion-Avoidance, Federation
Tables 22, 23, 25, 26

We updated VNode system to next version and we will show our recent progress of VNode system, especially focusing on applications working over VNode system. We are preparing four kinds of demos one for video multicasting and transcoding in virtual network, second for congestion avoidance method using virtual networks, third for FLARE switch, and at last for federation between different virtualization platforms.

In GEC20, video transcoding will be executed on both VNode and ProtoGENI side and comparison result will be shown, new virtual-network application (congestion avoidance) will be demonstrated, high-end FLARE network virtualization node with many 10G ports and three-way federation based on the common-API/slice definitionV2.0 will be demonstrated.

Developers interested in international/heterogeneous virtual network should see this demonstration, because it will show real execution results of VNode-ProtoGENI federation.

Participants:
* Univ. of Tokyo
  * Akihiro Nakao, nakao@iii.u-tokyo.ac.jp
  * Shu Yamamoto, shu@iii.u-tokyo.ac.jp
* Toshiaki Tarui, toshiaki.tarui.my@hitachi.com, Hitachi

GENI/FIRE Federation
Table 24

Demo description paragraph (three sentence minimum): The goal of this demo is to show how testbeds on both sides of the Atlantic can be federated. This demo will make use of a slice that stitches resources from GENI and FIRE testbed. !LabWiki will then be used to execute an experiment that makes use of the combined set of resources.

List of equipment that will need AC connections (e.g. laptop, switch, monitor): Just put in the number of connections needed if your demo description already lists equipment. 2 laptops

Total number of wired network connections (sum standard IP and VLAN connections): 2 standard IP Number of wired layer 2 VLANs (if any): Specify VLAN number, if known, approximate bandwidth, and whether tagged or untagged.

Participants:
* Mike Zink, zink@cs.umass.edu, University of Massachusetts

OpenBIDS
Table 27

Project description Project demo involves displaying the software (Network Intrusion detection system developed over OpenFlow). Performances are compared with SNORT and the statistics are studied in detail.

Participants:
* sxr1043@rit.edu, Rochester Institute of Technology

FIRE-GENI SDX 1.0
Table 28

We have created a first prototype of an SDX for exchange of stitching VLANs between European FIRE testbeds and GENI testbeds. We also developed a toolset to measure on certain moments in time the bandwidth performance on stitched links to see what an experimenter can expect and to learn more about the stitching technology.

Participants:
* Brecht Vermeulen, brecht.vermeulen@iminds.be, iMinds
* Rob Ricci, ricci@cs.utah.edu, Univ. of Utah
* Tom Lehman
* Xi Yang

Network Resilience with OpenFlow
Table 29

The University of Kansas and KanREN regional network will demonstrate network resilence using the KanREN/GENI and InstaGENI Openflow infrastructure.

Participants:
* James Sterbenz, jpgs@ittc.ku.edu, Univ. of Kansas

SDXs: Software Defined Networking Exchanges

Table 30

We will have a poster to show the current prototype for Software Define Networking Exchanges(SDXs) at !StarLight which will serve as exchange point for projects of members of the Consortium For International Advanced Networking Research using iGENI as a platform for Advanced Network Research and other Science Research.

Participants:
* Jim Chen, jim-chen@northwestern.edu, Northwestern Univ
* Joe Mambretti, j-mambretti@northwestern.edu, Northwestern Univ

Intelligent Data Movement System (IDMS)
Table 31

The IDMS project will demonstrate a data movement system running on GENI and highlight updates since GEC19. New areas include dynamically adding new resources to the service based on demand, stable images on both ExoGENI and InstaGENI, and improved monitoring of provisioned resources. In addition, we will show how the IDMS service can be used by other experimenters via shared vlans at available aggregates.

Participants:
* Ezra Kissel, kissel@cis.udel.edu, Univ. of Delaware

GENI Science Shakedown
Table 32

This demonstration will show recent progress made on the GENI Science Shakedown project. Topics include 1) techniques for importing ExoGENi images into InstaGENI and 2) running MPI applications on ExoGENI and InstaGENI. In addition, we will show an initial evaluation of both testbeds with respect to performance of MPI applications.

* Paul Ruth, pruth@renci.org, RENCI

ExoGENI
Table 32

Participants:
* Ilya Baldin, ibaldin@renci.org, RENCI
* Paul Ruth, pruth@renci.org, RENCI
* Jonathan Mills, jonmills@renci.org, RENCI
* Chris Heermann, ckh@renci.org, RENCI

SDX on GENI
Table 33

This demonstration will provide an update Software Defined Exchange (SDX) implementation work being led by the Georgia Tech team. The demonstration will provide an update on the work to deploy our SDX implementation on GENI. We will also update the progress toward an SDX API for GENI experimenters.

Participants:
* Russ Clark, russ.clark@gatech.edu, Georgia Tech

GENI Cinema
Table 34

An OpenFlow-based video hosting and streaming solution. Clients produce and consume video content to a video server behind an OpenFlow network. OpenFlow allows for path selection on the video server's private network.

* Ryan Izard, rizard@clemson.edu, Clemson Univ.
* Parmesh Ramanathan, parmesh@ece.wisc.edu, Univ. of Wisconsin-Madison
* Kuang-Ching Wang, kwang@clemson.edu, Clemson Univ.

QoS SDN
Table 34

We intend to demonstrate the power/necessity of QoS in our video streaming architecture. The required QoS can be enforced with the aid of SDN and openflow switches in a efficient manner. Our demonstration will consist of a video stream from a source such as a camera in a lecture room. This video stream will duplicated into two streams, both of which follow identical paths along the network along with the presence of various other traffic present. On one path, QoS will be deployed whereas on another path, the default service will be maintained and the difference will be observed!

Participants:
* Aditya Prakash, aprakash6@wisc.edu, University of Wisconsin-Madison

Posters

Measurements of Netmap-Based Data Plane Acceleration on xDPd
Poster

The poster will illustrate and report our current measurement results in characterizing the performance improvements achieved through netmap integration with xDPd. xDPd is an open source, modular software switch project:

https://www.codebasin.net/redmine/projects/xdpd/wiki. We recently started to collaborate with the FIRE team in Berlin, BISDN, on data plane acceleration for xDPd through netmap integration. Our measurements are in progress as xDPd also evolves. We also would like to receive GENI community feedback on such measurement challenges.

Participants:
* Levent Dane, ldane@uh.edu, Univ. of Houston
* Deniz Gurkan, dgurkan@central.uh.edu, Univ. of Houston

Integrating Dell SDP on GENI
Poster

The poster will present the state of the integration on GENI, our ideas and considerations on an aggregate manager for the Split Data Plane (SDP) switches on GENI. Dell has decided to donate ~10 SDP switches on GENI. UH is placing them on GENI rack locations, both ExoGENI and InstaGENI. The SDP switches should be considered GENI resources. However, there are challenges in sharing such a resource for simultaneous experimentation. We would like to present such considerations and receive discussions and feedback from GENI community.

Participants:
* Levent Dane, ldane@uh.edu, Univ. of Houston
* Deniz Gurkan, dgurkan@central.uh.edu, Univ. of Houston

End-to-End Cloud Security
Poster

End to End cloud Security is a project that scans the intermediate hosts for vulnerabilities between any cloud provider and the client. It performs scans on those nodes and collects and stores that data for further analysis. End to End cloud Security is a project that scans the intermediate hosts for vulnerabilities between any cloud provider and the client. It performs scans on those nodes and collects and stores that data for further analysis.

Participants:
* msm3850@rit.edu

Reliable Communications Protocol for Advanced Manufacturing
Poster

We will use Software Defined Networking (SDN) to deploy a protocol that can create and actively manage redundant paths between communicating devices and that will proactively build new paths before network congestion and faults degrade application performance. With replication of data packets and monitoring of multiple path performance, Internet communication will move beyond “best effort”

packet forwarding of TCP/IP to reliable packet delivery. Thus, the Reliable Communication Protocol (RCP) project will create a fundamentally new Internet capability.

Combining redundant paths with gigabit broadband will create an ideal platform on which to build any number of time- and reliability-sensitive applications for which today's Internet is insufficient or incapable. While our initial application focus is to support advanced manufacturing processes, there is clear potential to support applications in many other areas.

Participants:

* Alison Chan, chan7781@kettering.edu, Kettering Univ.