Enabling Campus Edge Computing using GENI Racks and Mobile Resources

Abhimanyu Gosain, Mark Berman,
Marshall Brinn, Thomas Mitchell
Raytheon BBN Technologies
Cambridge, USA
{agosain,mberman,mbrinn,tmitchell}@bbn.com

Chuan Li, Yuehua Wang, Hai Jin,
Jing Hua, Hongwei Zhang
Wayne State University
Detroit, USA
{chuan,ywang,haijin,jinghua,hongwei}@wayne.edu

Abstract—**This paper presents the architecture of GENI edge cloud computing network in the form of compute and storage resources, a mobile 4G cellular edge and a high speed campus network connecting these components. This deployment is available across fifty campuses in the US, all interconnected via a nationwide Layer-2 network. We present these capabilities in the context of vehicular sensing and control applications running on police patrol cars on the Wayne State University campus allowing end – users and researchers to collect rich datasets for public safety surveillance, vehicle internal-state sensing and modeling, and emulating next generation connected vehicle technologies. In particular, the paper provides insights about the usefulness of local edge computing cloud infrastructure for novel connected vehicle applications with high sensitivity to latency and bandwidth.**

Keywords—GENI Racks, Distributed Testbeds, Network Slicing, Mobile Edge Computing, Virtual Basestation, 4G small cell, Vehicular Networks, Connected Automated Vehicles (CAV)

I. INTRODUCTION : LOCATION MATTERS

It is difficult to overstate the importance of location when designing and operating distributed and networked computing systems, particularly those with demanding performance requirements. Application developers and service providers face a number of challenges when balancing advanced capability against performance needs. Not only do the physical locations of people, devices, vehicles, sensors, and instruments restrict the available design space, but current trends in cloud computing are exacerbating the problem by consolidating shared computing and storage resources in a small number of very large data centers. The challenges are particularly severe in latency-sensitive domains, including interactive and collaborative applications, real-time control, and safety critical applications. These systems often cannot afford the delays and variability of round trip times (RTT) to large data centers or the data movement implications of a centralized processing model.

It was no doubt in anticipation of this dilemma that a prescient scientist made a simple request: ``Give me the place to stand, and I shall move the Earth." [1] Future Internet and distributed cloud (FIDC) testbeds provide networked applications and services with the place to stand in the form of an edge computing capability that combines a distributed hardware footprint with deeply programmable networking, storage, and computing resources. One such FIDC testbed is GENI, the Global Environment for Network Innovations. [2]

GENI's edge computing strategy proceeds by deploying self-contained packages of network, computing, and storage resources, or *GENI Racks*, both at edge sites and at topologically interesting locations within the GENI network. The current GENI deployment includes over fifty such racks across the US as shown in Fig. 1. These racks are connected to each other and to a growing international federation of FIDC testbeds. Approximately one quarter of the current GENI sites also incorporate one or more GENI wireless 4G cellular base stations (BS). These wireless-capable sites provide a compelling wireless edge computing platform that incorporates single hop access to compute, storage, and network resources from vehicular and handheld mobile devices.

This paper presents GENI's edge computing capabilities and design, both in general and in the context of a specific mobile edge application in a public safety domain. We begin in Section II with a detailed view of the software services in GENI that allow for rapid

---

GENI is funded by the US National Science Foundation (NSF). Any opinions, findings, conclusions or recommendations expressed in this material are the authors' and do not necessarily reflect the views of the NSF.

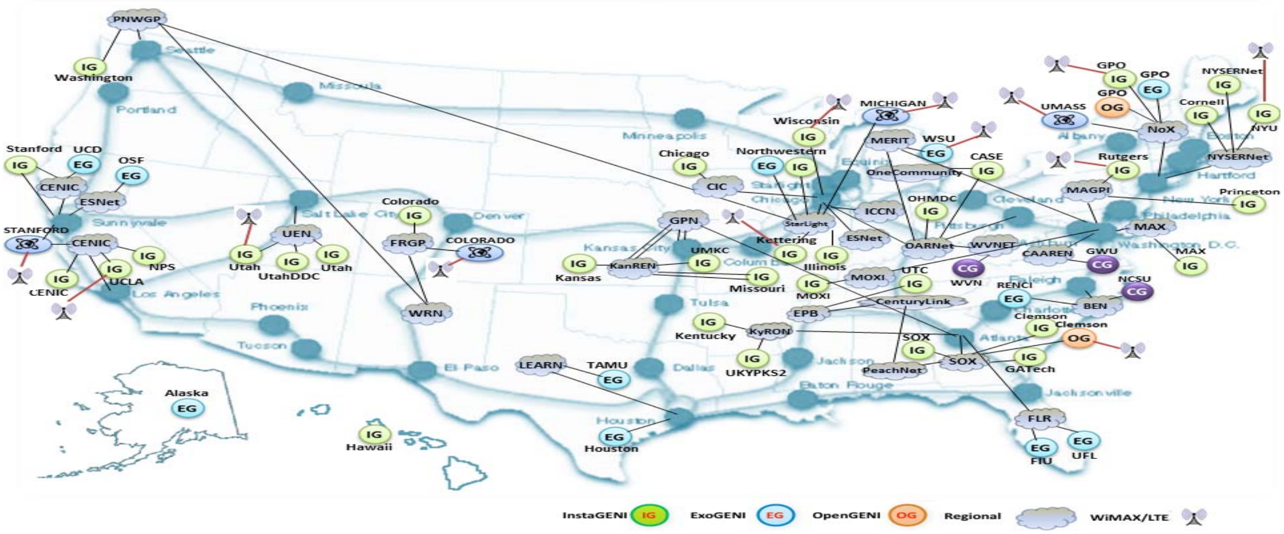

Figure 1: GENI Deployment Map

instantiation of an edge cloud environment and then describe the *GENI rack* internals. Section III outlines the last hop 4G cellular infrastructure and its integration with the edge cloud. Section IV outlines the work by researchers at Wayne State University in Detroit, MI working in collaboration with the university police department, who have deployed vehicle-mounted video cameras and vehicle internal state sensors in police cruisers policing their urban campus. The resulting data streams are continuously fed to the campus GENI rack, where they are processed in real-time and integrated into a large scale emulation to evaluate current and future CAV applications. Some related work is described in Section V and we summarize the main points and provide future directions in Section VI.

II. GENI RACKS AS FOUNDATION OF EDGE COMPUTING

GENI provides experimenters with isolated, deeply programmable topologies of virtual and physical resources in response to their requirements and requests, subject to resource availability called *slices*.

A *GENI slice* gives its owner control over some combination of virtualized substrate resources collected and assigned to the slice, which may include virtual servers, storage, programmable network elements, mobile/wireless platforms, and other programmable infrastructure components attached to the cloud network. *GENI slices* are built-to-order for the needs of each experiment.

Our ultimate goal is to manage the network and mobile substrate as a first-class resource that can be co-scheduled and co-allocated along with other resources, to instantiate a complete built- to-order edge network hosting a guest application, service, network experiment, or software environment.

Fundamental to this goal is the broad deployment of *GENI racks*, which enable a wide range of possibilities for rich edge computing applications and configurations.

A. Edge Computing at GENI Racks

GENI Racks consist of two switches, one for data plane Layer-2 (L2) traffic and one for control plane Layer-3 (L3) traffic. In addition, the racks provide compute/storage nodes from which full PC's or Virtual Machines (VM's) can be allocated on demand.

A software service known as the GENI Aggregate Manager (AM) manages access to the resources of a rack, providing these resources to authenticated, authorized users. The AM allocates computation and storage in the form of VM's and PC's. In addition, the AM allocates network resources in the form of isolated data-plane networks at L2 (pre-allocated VLAN tags, new L2 circuits). These networks may be OpenFlow [3] controlled. These racks open up a range of possibilities for edge computation applications. For example:

Rack-internal topologies. The rack allows for allocating a set of compute resources and making them connected on the data plane internal to that rack, providing rich network topologies internal to a single rack.

Campus Resources. The L2 switch can be connected to campus resources (other networks, compute resources, or sensors/actuators) on particular VLAN tags or subnets the connection to which can be managed by the AM. In this way, topologies may be constructed that contain a hybrid of GENI and campus-internal resources all interoperating at L2. For example, in the case study discussed in Section IV below, a local GENI rack is used to bridge between mobile devices and GPU-supported computing resources at a fixed campus site at Wayne State University.

Software Services. Users can establish public services on their allocated compute resources that can be available to users within or outside a given campus or domain.

Switch Management. The AM provides some degree of control over the L2 switch by enabling connections between ports, managing QoS requirements on particular connections, providing real-time data on port traffic and providing for establishing an OpenFlow controller for traffic on a particular network (independent of other traffic).

Cross-domain Stitching. The AM participates in a stitching protocol to enable topologies in one rack to connect at L2 to other topologies at other racks, potentially on other campuses or administrative domains.

Fig. 2 illustrates the range of edge computing options available from the vantage of a GENI rack.

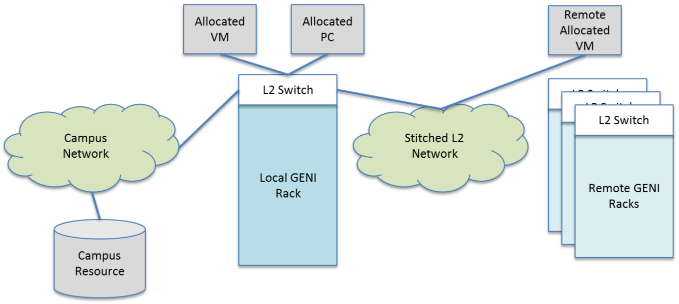

Figure 2: A campus-deployed GENI rack enables a range of edge-computing applications and topologies.

B. GENI Rack Internals

There are three fundamental flavors of GENI Racks: InstaGENI, ExoGENI [4] and OpenGENI [5] on which this paper will focus.

- InstaGENI is based on the Emulab [6] software suite developed by the University of Utah.
- ExoGENI is based on the ORCA software suite developed by RENCI.
- OpenGENI is based on the GRAM software suite developed by BBN.

ExoGENI and OpenGENI are, in turn, both built on the OpenStack [7] software stack, including quantum/neutron for network management and nova for compute resource management.

Each rack exposes a GENI AM through which it allocates and manages resources for authorized experimenters. Additionally, the AM manages the L2 switch to allocate and configure VLAN's and associate VLAN's and ports with traffic for specific compute resource interfaces.

Racks have a pre-allocated set of public-facing IP addresses that may be assigned to interfaces on allocated compute resources. In addition, the have a pre-allocated set of VLAN tags that may be assigned to interfaces on allocated compute resources. This management of VLAN tags is part of the stitching protocol that enables cross-campus isolated topologies, to be described below.

C. GENI Stitching: Getting to the Edge

The promise of edge computing is both the ability to perform rich computation, network and storage operations at the edge of the network, but to provide connectivity between these edges to allow for richer distributed applications.

GENI supports a protocol known as stitching that enables L2 connectivity between resources allocated on GENI racks. GENI racks are typically deployed on campuses connected to a L2 network (in addition to a L3 control plane network). The L2 network is typically connected to regional providers (e.g. CENIC, SOX) and then to national providers (e.g. Internet2). These internal providers provide services to allocate pre-provisioned VLAN tags or create new circuits on demand (using OESS [8], AL2S [9] or NSI [10], e.g.). They also provide VLAN translation services to allow traffic from one edge rack and traffic from another edge rack, with independently allocated VLAN tags, to interoperate.

Stitching is based on a protocol involving the edge rack AM's, the AM's at the internal networks and the GENI Stitching Computation Service (SCS). A GENI experimenter or user who wishes to allocate a cross-domain topology uses a tool such as stitcher.py or the

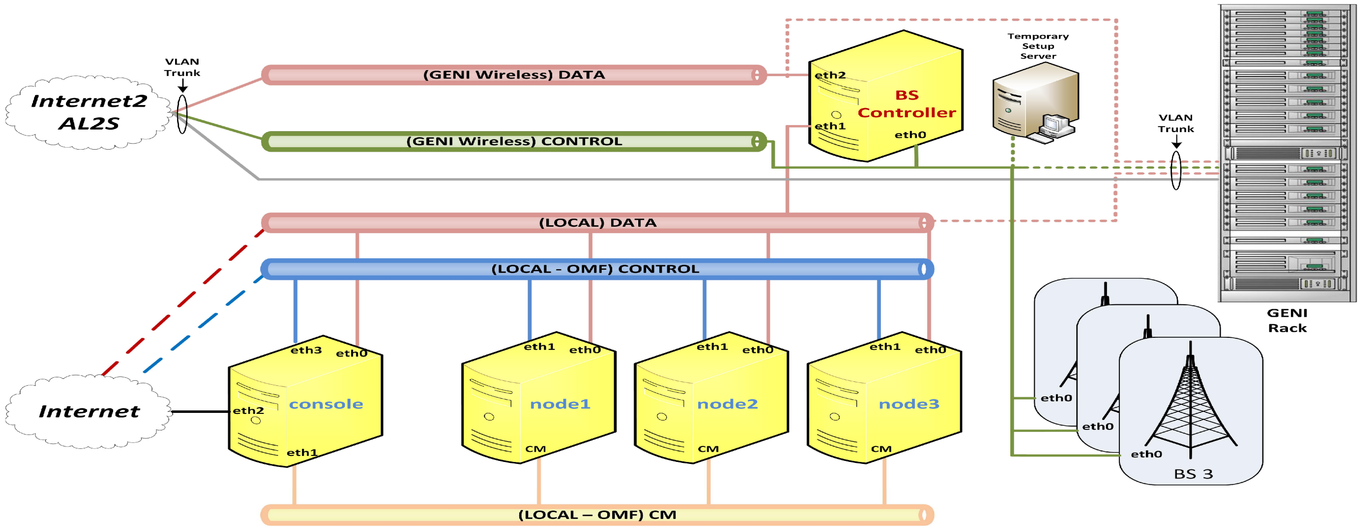

Figure 3: GENI Mobile Edge Campus Deployment

GENI Portal [11] or GENI Desktop [12] to specify the desired topology. The tool then consults the SCS to find the appropriate available L2 path between the specified edge resources. The tool then orchestrates the allocation of VLAN tags and translation sites and rules so that traffic will flow between (and only between) resources of the same allocated topology.

D. Securing GENI Edge Computing

The benefits of edge computing, including flexible, low-cost, highly controllable and configurable topologies, often come at the cost of security. Campus IT is often wary of allowing cloud resources to connect to campus resources without sufficient assurances about the reliability of these external parties.

GENI edge computing through GENI racks contains several aspects that help address these concerns:

- **Authentication**: GENI racks are federated with the GENI Clearinghouse and thus accept requests from users whose identity is trusted and vetted by the GENI federation.
- **Authorization**. GENI racks apply federation as well as local policy to ensure that no one is able to access more or different resources (compute or network or storage) than that to which they are entitled.
- **Accountability**. All GENI actions are logged and monitored so that a mis-behaving topology (intentionally or unintentionally) can be identified, shutdown and the person or people responsible for activity on these resources can be identified.
- **Isolation.** GENI manages VLAN allocation to assure L2 isolation between topologies of different experiments such that traffic between different experiments cannot be mixed, nor can any experiment access campus resources without appropriate authorization.

III. GENI MOBILE EDGE DEPLOYMENT

GENI wireless cellular base station deployments at a select number of campuses in the GENI testbed (shown in Fig. 1) provide the last hop mobile edge capabilities. The BS operates in the licensed Educational Broadband spectrum (EBS) and coexists with licensed service providers. Research described in this paper uses Airspan WiMAX BSs. The macro cell BS has Self organizing network (SON) support allowing for adaptive transmit power. A technology refresh deployment is underway to install LTE BSs at all GENI wireless sites. See Table 1 for comparison of radio technologies.

Table 1. Radio Interface Comparison

Parameter	*WiMAX*	*LTE*
Basestation Manufacturer	Airspan Air4G	Airspan AirHarmony
Cell Size	Macro	Femto,Pico
Channel	2.5-2.7GHz	Band 38,41 2.5-2.7GHz
Bandwidth	10MHz	20MHz
Max Tx Power	2 * 40dbm	37dbm

The BS controller software based on the Virtual BS system [13] extends the isolated data-plane networks at L2 to the mobile edge. The WiMAX BS is programmable through the ORBIT management

framework (OMF) [14], which manages GENI AM functionality for the wireless edge.
Fig. 3 shows the deployment of multiple small cell GENI Wireless BS(s) on a campus connected via high speed fiber backhaul to a BS controller.

The BS controller is co-located with the GENI rack and is connected via L2 to the data plane switch over two VLAN's. The *wimax-local* VLAN is point to point between the wireless edge, local GENI compute resources and campus resources. This extends the rack-internal topologies as described in Section II-A. The *wimax-multipoint* VLAN extends from the GENI rack to an internal provider all the way to Internet2, which performs L2 aggregation for each of the GENI wireless sites. This supports the cross-domain topologies mentioned in Section II-A. The net result of this implementation is the ability to reserve sliced configurations that connect wireless edge resources to a local GENI rack, remote racks, as well as local and remote campus resources, all without reliance on a commercial wireless carrier network or commodity internet service.

Mobile devices in form of Linux compatible USB-connected modems and Android OS handsets provide end user access to this edge cloud for compute intensive jobs and processing.

IV. CASE STUDY: SYMBIOTIC CAV APPLICATIONS AND EXPERIMENTS USING GENI EDGE COMPUTING

In this section, we elaborate on symbiotic connected-and-automated-vehicles (CAV) applications and experiments at Wayne State University (WSU) to demonstrate how GENI edge resources can be used by the application and research communities. In what follows, we first present the motivation for symbiotic CAV applications and experiments and describe the basic requirements imposed by such applications and experiments; then, we provide details about the designs and implementations of the example applications and experiments.

A. Motivation

Vehicle accidents cause over 1.4 million fatalities and 50 million injuries per year across the world, and motor vehicles account for over 20% of the world's energy use and over 60% of the world's ozone pollution. Transforming the traditional, single-vehicle-based safety and efficiency control, next-generation vehicles will cooperate with one another and with the transportation infrastructures (e.g., traffic lights) to improve transportation safety and efficiency [15]. For instance, vehicles can coordinate with one another to avoid collisions; based on real-time road and traffic conditions as well as the control actions of surrounding vehicles, a vehicle can control its speed, throttle, and gear to improve fuel economy by as much as 50% or more in certain situations. [16,17]. Besides its transformative impact on transportation, CAVs can also serve as mobile sensing platforms and contribute to smart cities and communities. Complementing infrastructure-mounted cameras, for instance, vehicle-mounted cameras can provide mobile surveillance information for public safety surveillance and emergency response.

Even though CAVs have been extensively studied, we are still at the infancy of CAV development and deployment, and there exists a wide spectrum of technological and application challenges before we can fully realize the potentials of CAV. These challenges, together with the associated societal challenges, require us to rethink innovation paradigms to speed up the research, development, and deployment of CAV networks and applications. Leveraging GENI platforms and infrastructures for CAV networking and application deployment, one promising approach is to allow symbiotic exploration and evolution of CAV technologies and applications in shared real-world systems and environments.

B. Enabling CAV applications and experiments in GENI

Two distinct CAV applications (i.e., 3D mapping for public safety surveillance and vehicle internal-state sensing) and a CAV emulation experiment are selected in this section to demonstrate the capability of GENI in supporting symbiotic CAV applications and experiments. Specifically, in 3D mapping for public safety surveillance, a 3D map of WSU campus is built with rich context information for interpreting real-time video surveillance views. To enable vehicle internal-state sensing and analytics, multiple sensors are equipped, and a large amount of real-time sensing data need to be archived in GENI racks. The CAV emulation experiment is an at-scale, high-fidelity experiment to evaluate current and future CAV wireless networking solutions. These applications and experiment can all benefit from high-speed computing, high-volume storage, as well as high-speed, low-latency communication at network edges, as exemplified by GENI edge computing and networking facilities.

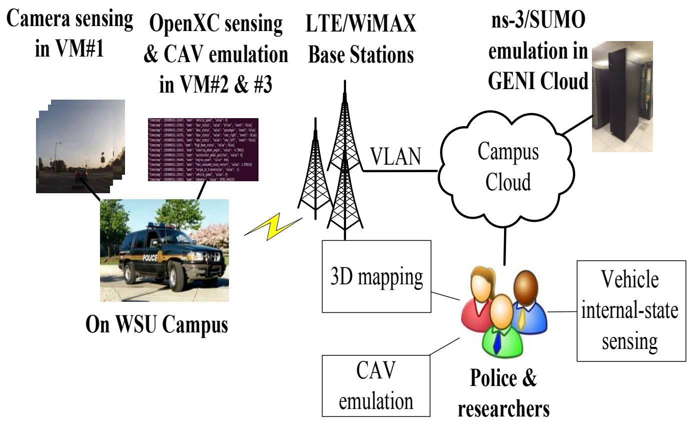

Figure 4: Symbiotic CAV applications and experiments using GENI

Fig. 4 depicts a typical scenario of the aforementioned CAV applications and experiments using GENI at WSU. It involves vehicular sensing and control (VSC) platforms, LTE/WiMAX base stations, GENI edge computing infrastructures (i.e., GENI racks), and users. The VSC platforms deployed on WSU police patrol vehicles enable high-fidelity vehicular sensing, computing, and wireless networking. Through virtualization, the vehicular sensing, computing, and networking resources of the VSC platforms are partitioned into multiple network slices such that different slices are allocated to different CAV applications and experiments [18]. Through the GENI WiMAX/LTE network on WSU campus and GENI VLAN within WSU campus (denoted as WSU VLAN), the VSC platform slices are integrated with GENI edge computing infrastructures and other campus infrastructures (e.g., computing and networking resources at researchers' labs) to enable the exploration of novel CAV applications and technologies. CAV applications and experiments leverage high-fidelity data from the VSC platforms through the GENI mobile edge infrastructure.

(i) 3D mapping for public safety surveillance

We create an end-to-end network slice spanning the VSC platforms on WSU campus, GENI WiMAX/LTE network, WSU VLAN, as well as end-user and researcher infrastructures. Within the slice, dedicated communication and computing resources are reserved for streaming real-time video streams from field-deployed vehicles to police control centers. To provide a global, naturalistic view of the campus and to facilitate public safety surveillance, a 3D mapping application is deployed on the GENI rack at WSU, and it fuses all the camera sensing data from the infrastructure cameras and VSC platforms to create an augmented 3D world for public safety officers to navigate through. With the 3D mapping application, an officer can navigate through the campus as if he/she can examine the campus from arbitrary view-points and view-angles without having to examine tends of screens separately as in existing public safety surveillance systems.

Fig. 5 shows an example of positioning real-time camera views from WSU police patrol vehicles in the 3D augmented reality environment. Fig. 6 provides an example of embedding real-time camera views into the 3D campus virtual reality environment so that the officers in the control center can examine field just as if they were in field (instead of in the control center).

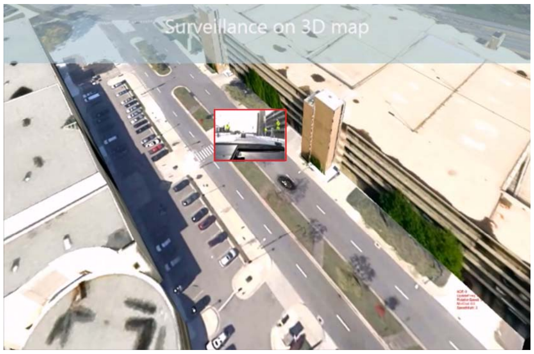

Figure 5: Surveillance in a 3D map

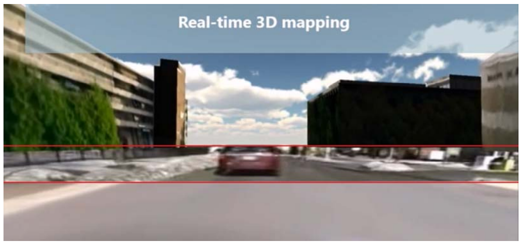

Figure 6: Real-time 3D mapping

For the 3D mapping application, it is important to be able to process the large volume of camera sensing data close to the network edge so as to avoid the extra communication bandwidth and delay between the network edge and remotely-located central cloud infrastructures. This reduction in communication delay is beneficial to real-time 3D mapping, especially considering the fact that the edge wireless communication may already introduce non-negligible delay during inherent uncertainties in wireless communication itself. For instance, Fig. 7 shows the round-trip time (RTT) between the VSC platforms and WSU GENI rack, as a function of the distance between the WSU GENI rack and the vehicles carrying the VSC platforms. We see that, due to spatiotemporal variations of wireless channels (e.g., as a result of different environmental conditions such as surrounding buildings in different directions), there is a significant variation in RTT for a given distance between the GENI rack and VSC platforms, which makes the worst-case delay non-negligible.

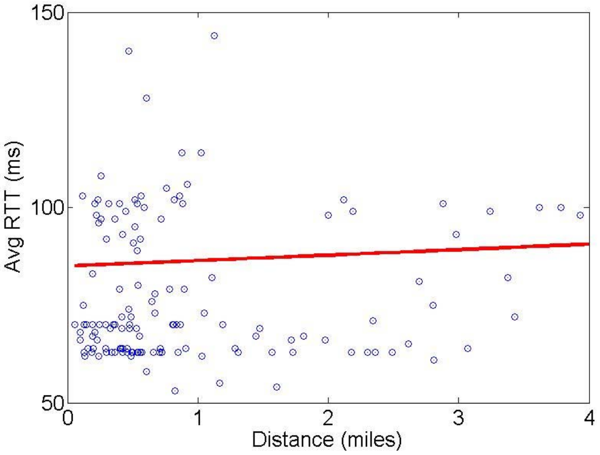

Figure 7: RTT between VSC platforms and WSU GENI rack

(ii) Vehicle internal-state sensing

Using the same police vehicle deployment as shown in Fig. 4, another slice is created for sensing vehicle internal state via OpenXC [19], for purposes such as diagnostics, prognostics, and fuel economy sensing and optimization [15]. In this application, a large amount of data needs to be stored so that they can be analyzed for identifying fuel economy optimization methods. Additionally, vehicle internal state (e.g., those related to driving behavior such as acceleration) is needed to enable CAV emulation which is discussed in Section IV-B (iii), and this needs to be communicated in real-time. To this end, high speed WiMAX/LTE links and high performance edge racks are reserved for this slice. Depending on the model of vehicles, the VSC platforms provide users with more than 10 types of vehicle state such as speed, engine speed, fuel-level, odometers, gear, and torque. Besides enabling real-time CAV emulation to be discussed shortly, this sensing data can be stored for offline analysis. For instance, users can observe this data visually before conducting rigorous analytics. Fig. 8 shows a snapshot of vehicle speed, torque, and engine speed in the visual gauges, and Fig. 9 shows the vehicle fuel economy, speed, fuel consumption, and odometer values. This data can be very useful for studying vehicles' properties and drivers' behaviors. Storing the data at the network edge not only helps reduce communication delay which is undesirable for real-time CAV emulation, it also helps reduce the requirement for backbone communication bandwidth.

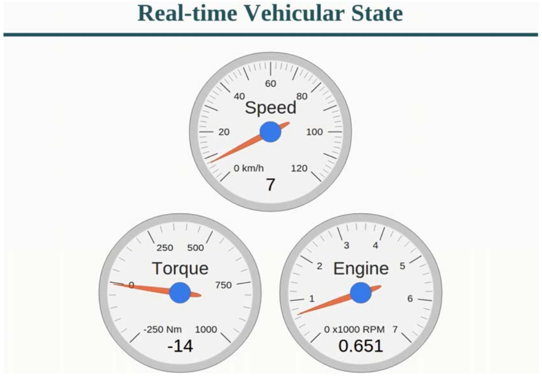

Figure 8: Real-time vehicle state sensing

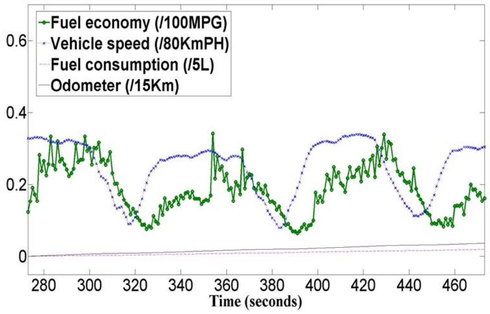

Figure 9: Fuel economy sensing

(iii) CAV emulation

A third network slice is created for CAV emulation, which integrates vehicles in real-world traffic with multi-domain simulation of V2X communication, vehicle dynamics, and traffic flow in high-performance edge computing infrastructures of GENI, thus enabling multi-domain emulation (MDE) where the high-fidelity of real-world vehicle traffic is integrated with the flexibility and scalability of simulation in the GENI edge computing infrastructures [18]. For instance, using the OpenXC field sensing data, we can instantiate the microscopic vehicle mobility models of SUMO to enhance the fidelity of traffic dynamics simulation, an example being the instantiation of a key model parameter "driver reaction time" as show in Fig. 9. Similar to the network slices reserved for the two CAV applications discussed earlier, dedicated communication and computing resources are reserved for the CAV emulation, including dedicated communication resources from the VSC platforms to the WSU GENI rack.

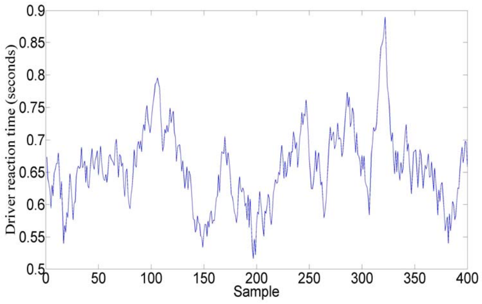

Figure 9: Example model instantiation: driver reaction time

With the instantiated high-fidelity simulation models in our CAV emulation system, we can comparatively evaluate different VSC networking and application solutions. For instance, using information conveyed by inter-vehicle wireless communication, a vehicle can estimate inter-vehicle distances between those vehicles beyond the line-of-sight (LOS) sensing ranges of radars, and the reliability of inter-vehicle communication significantly impacts the estimation accuracy since inter-vehicle communication reliability directly impacts the performance for a vehicle to track the locations of vehicles beyond LOS. For illustration purpose, we consider two inter-vehicle communication protocols for medium access control: *CSMA* which uses carrier sensing multiple access method to control wireless channel access by individual vehicles, and *vPRKS* which controls channel access based the wireless transmission scheduling protocol PRKS [15] which controls co-channel wireless interference in a predictable manner and thus enables predictable inter-vehicle communication reliability.

In our emulation experiment, we consider a real-world road topology representing a 2×2 km^2 area of the WSU campus in Detroit, USA. To enable high-fidelity emulation of CAV networks, traffic infrastructures and rules of road usage such as traffic lights, speed limits, one-way roads, and right-of-way rules are considered, and they allow us to simulate common vehicular situations such as overtakes and stops at intersections. In our CAV emulation, a few real-world vehicles (e.g., those with VSC platforms on WSU campus) are used while the rest of the vehicles are emulated in the WSU GENI rack. The real-world vehicles keep injecting high-fidelity sensing and wireless-channel measurement data as well as application traffic into the emulation system in the WSU GENI rack. The emulated network is actually a hybrid network, where the real-world and emulated vehicles communicate and interact with each other using CAV protocols. Instead of large-scale real-world CAV testbeds which are costly and difficult to manage and operate, such CAV emulators enable experimenters to create target CAV scenarios and quickly evaluate current and future CAV technologies and applications in a low-cost and flexible manner.

For one emulation experiment, Fig. 10 shows the estimation errors using the two MAC protocols CSMA and vPRKS. Without predictable interference control of co-channel interference, communication reliability tends to be low in CSMA, especially when data traffic load is high (e.g., one packet every 50ms from every vehicle). With predictable interference control enabled by protocol PRKS [15], however, vPRKS ensures high communication reliability. Accordingly, the distance estimation error tends to be much lower in vPRKS as compared with that in CSMA, as shown in Fig. 10.

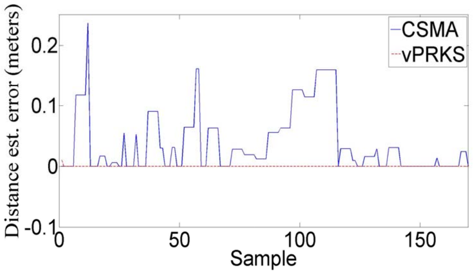

Figure 10: Example CAV solution evaluation: inter-vehicle distance estimation error

Similar to enabling the aforementioned emulation experiment on comparing different CAV communication protocols, the VSC platforms, WiMAX/LTE wireless networks, high-speed VLANs, and high-performance edge computing infrastructures of GENI will enable high-fidelity, at-scale evaluation of other CAV technologies and applications, thus helping develop the technology and application foundations for next-generation road transportation.

V. RELATED WORK

Previous testbed virtualization efforts and cloud computing infrastructures such as Planetlab [20] and Emulab have provided the central cloud functionality but with the added WAN delay and jitter. Satyanarayanan et al [21] have proposed a vision of using wireless LAN connectivity from mobile devices to leverage compute cycles and storage resources at co-located edge cloud infrastructure. PhantomNet [22] is a mobile networking testbed at University of Utah that provides researchers with a set of hardware and software resources that they can use to develop, debug, and evaluate their mobility ideas but in a lab environment. The international community also has various mobile computing testbeds under the Fed4Fire [23] project. NITOS [24] and VirtualWall [25] offer wireless testbed experimentation on mobile access technologies and emulated cloud topologies respectively. However, both initiatives fall short on providing an end to end at-scale environment for mobile edge computing.

By building GENI AM software that allows for rapid instantiation of edge cloud VM's and an isolated data plane connection from the mobile devices to the GENI rack VM's via a 4G connection, GENI provides an end to end ecosystem for interactive mobile cloud applications.

VI. SUMMARY AND FUTURE WORK

In this paper, we have presented the design and architecture of a distributed nationwide mobile edge computing infrastructure testbed; GENI. We have described the resources of this testbed that make it an attractive candidate for edge computing research and applications. We have shown this in the context of multiple computationally intensive CAV applications and experiments that collect and process high fidelity data streams using edge computing on GENI racks from patrol vehicles driving around a campus. Our intention is to share with readers how GENI edge infrastructure can enable symbiotic CAV applications and experiments. It is, however, not our intention to strive for the optimality of the approaches we have taken in developing the CAV applications and experiments, which would be a part of our future work.

Future work involves a strategic upgrade of the Cellular BS to LTE radios and a wider footprint of GENI racks to new campuses, and expanded interoperation with other compatible testbeds. We plan to harden the GENI testbed to allow for easier setup of edge computing applications. This will allow these applications to perform gracefully and deliver critical information for analysis faster and more efficiently. This will be accomplished by doing a refresh of our testbed hardware, which is underway. The mobility aspect of this testbed will meet the application requirements by implementing network slice isolation in the Evolved Packet Core (EPC) of an LTE network and also upgrade the GENI BS controller API to expose mobility parameters to researchers allowing for finer grained and real time control in their applications.

REFERENCES

[1] Pappus of Alexandria, Synagoge, Book VII, c. 340. (Quoting Archimedes, c. 287 BC - c. 212 BC.)

[2] Mark Berman, Jeffrey S. Chase, Lawrence Landweber, Akihiro Nakao, Max Ott, Dipankar Raychaudhuri, Robert Ricci, Ivan Seskar, GENI: A federated testbed for innovative network experiments, Computer Networks, Volume 61, 14 March 2014, Pages 5-23, ISSN 1389-1286, http://dx.doi.org/10.1016/j.bjp.2013.12.037.

[3] OpenFlow:http://archive.openflow.org/documents/openflow-wp-latest.pdf

[4] ExoGENI: http://www.exogeni.net [Online].

[5] OpenGENI: http://www.opengeni.net [Online]

[6] Brian White, Jay Lepreau, Leigh Stoller, Robert Ricci, Shashi Guruprasad, Mac Newbold, Mike Hibler, Chad Barb, and Abhijeet Joglekar, "An integrated experimental environment for distributed systems and networks," SIGOPS Operating Systems Review, vol. 36, no. SI, pp. 255-270, December 2002.

[7] OpenStack project website. Available: http://www.openstack.org

[8] OESS: https://globalnoc.iu.edu/sdn/oess.html [Online]

[9]AL2S:https://noc.net.internet2.edu/i2network/advanced-layer-2-service.html [Online]

[10]NSI:https://www.researchgate.net/publication/289207695_Network_service_interface_-_Gateway_for_future_network_services [Online]

[11] GENI portal, https://portal.geni.net/.

[12] GENI Desktop, https://genidesktop.netlab.uky.edu/

[13] Gautam Bhanage, Ivan Seskar, Rajesh Mahindra, Dipankar Raychaudhuri, "Virtual Basestation: Architecture for an Open Shared WiMAX Framework", ACM SIGCOMM VISA Workshop, 2010

[14] Thierry Rakotoarivelo, Max Ott, Guillaume Jourjon, Ivan Seskar, "OMF: a control and management framework for networking testbeds", in ACM SIGOPS Operating Systems Review 43 (4), 54-59, Jan. 2010.

[15] R. Johri, J. Rao, H. Yu, and H. Zhang, "A Multi-Scale Spatiotemporal Perspective of Connected and Automated Vehicles: Applications and Wireless Networking," *IEEE Intelligent Transportation Systems*, 8(2), 2016

[16] H. Zhang, X. Liu, C. Li, Y. Chen, X. Che, F. Lin, L. Y. Wang, and G. Yin, "Scheduling with Predictable Link Reliability for Wireless Networked Control," in IEEE/ACM IWQoS, 2015.

[17] H. Zhang, X. Che, X. Liu, and X. Ju, "Adaptive Instantiation of the Protocol Interference Model in Wireless Networked Sensing and Control," *ACM Transactions on Sensor Networks,* 10(2), 2014.

[18] Y. Wang, H. Jin, C. Li, H. Zhang, J. Hua, J. Rao, G. Riley, A. Holt, and P. Gossman, "Symbiotic CAV Evolution: Software-Defined Infrastructure and Case Study in Public Safety (working paper)," 2016.

[19] "OpenXC platform," http://openxcplatform.com/.

[20] L. Peterson, S. Muir, T. Roscoe, and A. Klingaman. PlanetLab Architecture: An Overview. Technical Report PDN–06–031, May 2006.

[21] Mahadev Satyanarayanan; Bahl, P.; Caceres, R; Davies, N., "The Case for VM-Based Cloudlets in Mobile Computing," in Pervasive Computing, IEEE , vol.8, no.4, pp.14-23, Oct.-Dec. 2009

[22] A. Banerjee, J. Cho, E. Eide, J. Duerig, B. Nguyen, R. Ricci, J. Van der Merwe, K. Webb, and G. Wong. PhantomNet: Research Infrastructure for Mobile Networking, Cloud Computing and Software-Defined Networking.

GetMobile: Mobile Computing and Communications. Volume 19 Issue 2, pp. 28-33.April 2015

[23] Wim Vandenberghe, Brecht Vermeulen, Piet Demeester, Alexander Willner, Symeon Papavassiliou, Anastasius Gavras, Michael Sioutis, Alina Quereilhac, Yahya Al-Hazmi, Felicia Lobillo, Florian Schreiner, Celia Velayos, Albert Vico-Oton, Georgios Androulidakis, Chrysa Papagianni, Okung Ntofon, and Michael Boniface, "Architecture for the Heterogeneous Federation of Future Internet Experimentation Facilities," in *Future Network and Mobile Summit 2013 Conference Proceedings*, 2013.

[24] VirtualWall: http://www.fed4fire.eu/virtual-wall/ [Online]

[25] K. Pechlivanidou, K. Katsalis, I. Igoumenos, D Katsaros, T. Korakis, L. Tassiulas. "NITOS Testbed: A Cloud based Wireless Experimentation Facility," *FIDC, 26th International Teletraffic Congress (ITC 26)*, 2014.