Evaluating 5G Multihoming Services in the MobilityFirst Future Internet Architecture

Parishad Karimi, Michael Sherman,
Ivan Seskar, Dipankar Raychaudhuri
WINLAB,
Rutgers University
{*parishad,msherman,seskar,ray*}*@winlab.rutgers.edu*

Francesco Bronzino
Inria, Paris
francesco.bronzino@inria.fr
Abhimanyu Gosain
Raytheon BBN Technologies
agosain@bbn.com

***Abstract*—In the recent years it has become increasingly evident that the current end-to-end host-centric communication paradigm will not be capable of meeting the ongoing demand for massive data rates and ultra-low latency. With the advent of fifth generation of cellular architecture (5G) to support these requirements on the wireless edge of the network, the need for core network solutions to play a complementary role is conspicuous. In this paper we present and tackle some of the challenges of deploying a Future Internet Architecture (FIA), called MobilityFirst (MF), specifically for 5G use case scenarios. We report our findings of the deployment based on a setup on a small-scale testbed (ORBIT) and a nation-wide distributed testbed (GENI), and illustrate some results for the use case of device multihoming, in comparison with current TCP/IP based solution, i.e. Multipath TCP.**

I. INTRODUCTION

The Internet protocol (TCP/IP) is currently coping with some of its fundamental limitations by deploying architectural fixes (mobility, security, multicasting, NAT, etc.) affixing themselves to a fixed architecture - which may serve a valuable short-term purpose, but significantly impairs the long-term flexibility, reliability, and manageability of the Internet. The FIA (Future Internet Architecture) program is just one of the steps towards the development of 5G networks.

The challenge of deploying and evaluating FIA models for 5G requires testbeds that have the inherent capacity, capability, reliability, programmability, availability and security to provide this infrastructure. The evaluation of MobilityFirst (as one of the leading projects of FIA) requires visibility into and control of the mobile core network and UE. Both the GENI (Global Environment for Network Innovations) wireless testbed [1] and ORBIT testbed [2] meet the criteria described above. GENI provides access to a fully programmable end-to-end 4G-LTE deployment at twelve campuses across the USA. Moreover, the ORBIT testbed maintains a variety of commercial and open source LTE infrastructure, both on commercial and on generic hardware, which support different potential configurations for UE (User Equipment), eNB, and EPC (Evolved Packet Core).

Having these testbeds available, a diverse set of effective use cases consisting of mobility and multicast services, content retrieval and publisher-subscriber systems, virtual networking, etc can be deployed and evaluated. In this paper we focus on the implementation of mobility services, namely device multihoming and its deployment in testbeds including 5G technologies. In Sec. III the principal design elements of the Mobilityfirst architecture are described along with how Mobilityfirst enables 5G services (specifically multihoming). Next in Sec. IV the details for the components employed in the implementation are elaborated on. In Sec. V the experimental methodology and results are presented. We discuss large scale implementation of the experimental setup in Sec. VI and finally Sec. VII concludes the paper.

II. RELATED WORK

One of the important issues regarding clean-slate Future Internet Architecture projects [3] is their deploy-ability on testbeds and in real-world scale networks. The clean-slate nature of these architecture proposals necessitates the study of approaches to deploy and evaluate them on todays networks along with 4G/5G technologies available on testbed and real-world scale. Previous work on deploying different experiments, ranges from virtualizing base stations [4] to services like network traffic shaping [5] on top of 4G technologies in ORBIT and GENI testbeds [6]. In [7] specific scenarios like routing and name resolution scalability in Mobilityfirst FIA architecture along with real-world experiments were conducted on ORBIT and GENI testbeds.

Despite its significance, combining the scope of FIA proposals with the advancements in wireless edge network (software-defined wireless networking and 5G technology) is a less explored area. In [8], authors describe the framework they implemented within GENI to test vertical handover between heterogeneous wireless networks. Our work in this paper focuses on deploying more general services within MobilityFirst architecture in the context of 5G technologies available within ORBIT testbed.

Deploying multi-interface connectivity has also been the subject of study in a couple of projects [9], [10]. The most-widely adopted multihoming technique used in practice is Multipath TCP [11]. There have been efforts focusing on micro-scale deployment of MPTCP and its interaction with diverse wireless access technologies like WiFi, LTE, etc.

[12]–[14]. In [15], a control and measurement SDN-based framework has been built on top of two testbeds to enable large-scale deployment, testing and evaluation of MPTCP.

III. MOBILITYFIRST OVERVIEW

MobilityFirst (MF) [16] is a name-based architecture founded based on the concept of “named objects” represented by flat globally unique identifiers (GUIDs). The data communication in the network will consist of self-defining and self-certifying packets which carry the source and destinations information attributes (embedded into the GUID identifier) and service intent (defined by an additional field called the SID) and can thus be processed in a very general manner by store-and-forward routers along the path. As a result, MF paves the way for transitioning today’s IP address-based communication to name-based communication which opens up opportunities for many services like multihoming [17], [18], multicast [19], anycast and DTN delivery. The enabling component which provides the mapping between GUID and network address(es) associated with that GUID is the Global Name Resolution Service (GNRS). Each network attached object will insert its *GUID-to-NA mapping* into GNRS, and update its entry with its current address(es) when changing points of attachment to the network.

Data delivery in Mobilityfirst is done based on Hop protocol [20], which involves transferring segmented chunks and retransmission among storage-aware routers. In contrast with the current Internet approach in TCP/IP which controls the transmission rate on an end-to-end window-controlled manner, Mobilityfirst incorporates hop-by-hop block transport protocol. The congestion control in this transport protocol follows a segment-level back pressure mechanism as follows:

Each router will send a control message CSYN to its next hop (after transmission of its current chunk, and prior to the next chunk). The next-hop router will respond with a control message C-ACK, containing bitmap information of the received packets. The router will proceed with the next chunk transmission only on receipt of a successful whole chunk delivery.

MobilityFirst enabling 5G services: Supporting mobile data services in wireless edge networks is one of the most significant use cases in Mobilityfirst, due to its built-in support for essential cellular network features including authentication and dynamic mobility, both micro-level handoff and macro-level roaming. This alleviates the need for conventional mobility gateways in both the control and data paths, thus eliminating tunneling overheads, reducing latency and avoiding traffic bottlenecks at the service gateway. Consequently, in a flat MF-based cellular network routers, base stations and access points run the standard MF protocol to realize mobility services in a fully decentralized manner with each radio access technology terminated at the AP/BS and simply plugged in to the network. The MF network can be used by cellular network operators seeking to improve efficiency and performance, as well as by an ISP aiming to introduce mobility services across heterogeneous access networks.

Multihoming in MobilityFirst: Providing hop-by-hop reliable delivery and name-address separation, Mobilityfirst facilitates inherent support of mobility and multihoming (in addition to many other services, which are not in the scope of this paper). As previously discussed, each multihomed device will possess a long-lasting GUID and a set of dynamic short-lasting network addresses (NAs). The MF name-based API offer basic messaging primitives like *send(GUID,data)* and *get(GUID,data)*. As the data traverses the routers within the network, they perform GNRS lookup and append the various network addresses associated with the destination GUID as needed to the chunks. The in-network entities will detect multiple interfaces on the end host through these chunks and a bifurcation point will schedule the data towards each interface based on fine-grained link quality information. An overview of the multihoming in Mobilityfirst along with the main components and functions of the architecture are illustrated in Fig.1

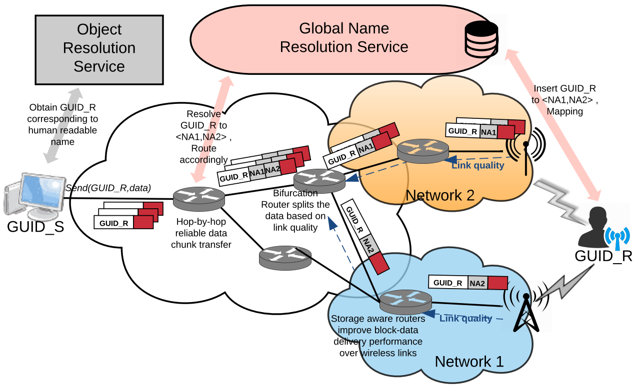

Fig. 1. Multihoming overview in Mobilityfirst Architecture. The baseline example shows how the server with GUID_S will transmit data to the client with GUID_R

IV. TESTBED IMPLEMENTATION

A. MF protocol stack

In order to move towards testbed based experimentation a prototype of the MobilityFirst architecture has been developed. The prototype - first described in full here [21] - includes the main components that are part of the designed architecture. The result of these efforts consisted of three main tools: a GNRS implementation based on DMap’s design [22], a Click [23] based software router that implements the GSTAR [24] routing protocol and a multiplatform protocol stack and network API for clients [25]. Applications and network services can be implemented as extensions of these basic elements. Moreover, the necessary support to automate experimentation were developed using the OMF [26] framework and provide statistic collections through OML [27]. While we refer to [21] for a complete description of the prototype, for the interest of this paper we focus on two core aspects that characterize the results: the multihoming mechanisms and how the two main components, i.e. the clients' stack and the routers, interface with the underlying network protocols.

Host Stack and API. The host network protocol stack has been implemented on Linux and Android platforms as a user-level process built as an event-based data pipeline. It implements the name-based network protocol primitives through the GUID service layer. A new network API [25] is available to applications to perform name based operations. Through per operation based options, application can request custom delivery service types, including multicast, anycast and multihoming. A policy-driven interface manager handles concurrency across available multiple interfaces. The device-level policies allow the user to manage how data is multiplexed across one or more active interfaces. Following the spirit of flexible deployability on top of multiple experimental scenarios, the stack has been enabled with an interface abstraction that can smartly adapt to different networking environments. These interchangeable *Interface* classes (figure 2), support deployment on top of different layers of the architecture, including: a) native support of the MobilityFirst protocols on top of a L2 network and overlay support both on top of b) barebone IP network or c) a full overlay solution on top of UDP. Thanks to this support for a wide range of overlay modes, it was possible to deploy the components on top of different access and network technologies, such as the GENI LTE infrastructure.

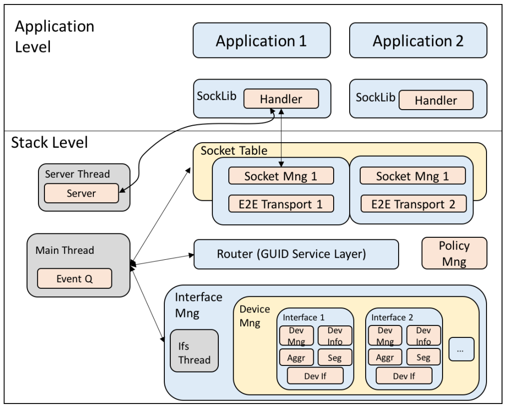

Fig. 2. Client host stack block diagram.

Routers. Software routers are implemented as a set of routing and forwarding elements within the Click [23] modular router. The router implements dynamic-binding using GNRS, hop-by-hop transport, and storage-aware routing. It integrates a large storage, an in-memory *hold buffer*, to temporarily hold data blocks when destination endpoints during short-lived disconnections or poor access connections. A particular instance of this system, implements what we call a MobilityFirst access router, a router providing access connectivity to clients. The router uses a similar interface abstraction to the one presented for the host stack to adapt to different environment conditions.

B. LTE SDR deployment

The LTE Testbed setup for evaluation consists of a Netgear 341U dongle with a custom SIM, a commercial base station (Airspan Airsynergy 2000) running in the ORBIT Grid, and an Amarisoft EPC instance running on the GENI Wireless VLAN. The eNB and client run on a custom built ORBIT nodes, with the MF software installed on them.

C. Stitching MF and LTE together

As MF routers need a L2 connection between them, we connect the LTE Client, and LTE MF router via an L2TP tunnel. This enables us to swap out the underlying LTE implementation transparently to the MF configuration. In addition to the tested configuration, we have tried a commercial eNB (Airspan Airsynergy 2000), as well as a soft UE and eNB setup (OpenAirInterface: noS1) [28], with no EPC. The shown configuration strikes a balance between customization and stability. A sample result of ping for 2 cases of LTE over IP vs LTE over MF is shown in 3. Mfping is the equivalent ping application modified to use mfAPI. As can be seen the ping values show comparable results and the slightly higher values for MF can be a result of using software routers and larger packets (extra name and SID headers). However as will be shown in Sec.V, this does not influence the data throughput performance for MF.

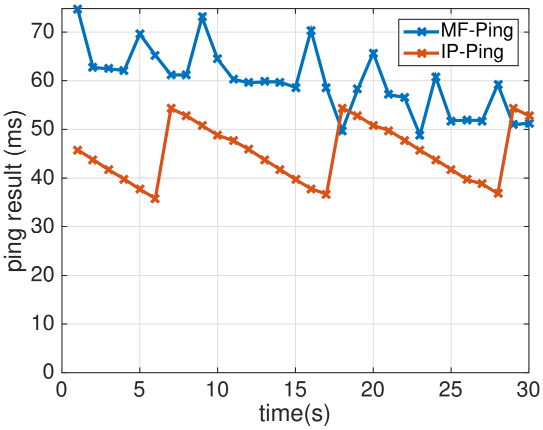

Fig. 3. Result of ping over the LTE link, for both IP and MF-based applications

V. EXPERIMENTAL EVALUATION

As discussed before, one of the important use cases in 5G architecture is device multihoming, where a user device is capable of connecting to multiple networks simultaneously. Multipath TCP is one of the widely-adopted solutions for device multihoming, where striping of data is handled by the transport layer, on an end-to-end basis. In contrast, Mobilityfirst pushes splitting of the data towards multiple interfaces to network elements, i.e. the *bifurcation* router. Through experimentation, we focus on evaluating how utilizing network elements for scheduling the traffic on each path towards

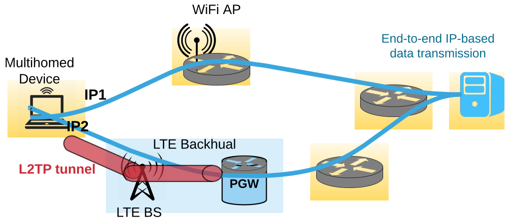

(a) MPTCP-based setup: data splitting on an end-to-end basis

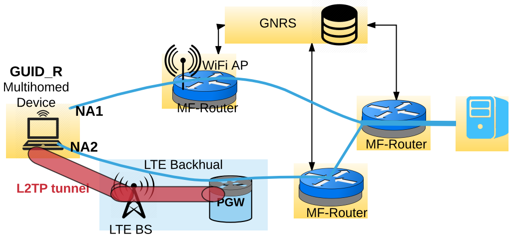

(b) MobilityFirst-based setup: GNRS service and MF-software routers enable network assisted multihoming

Fig. 4. Benchmark experimental setup

various interfaces for a 5G device, along with reliable hop-by-hop data delivery transport mechanism will perform in comparison with delegating scheduling and data-splitting to the end hosts in MPTCP.

A. *ORBIT*

In our experimental evaluation we aim to deploy a baseline device multihoming scenario in ORBIT testbed, in which the client will receive data on its both interfaces, Wi-Fi and LTE. The LTE backhaul is implemented as discussed in Sec.IV-C. The Wi-Fi link is using 802.11g technology (Atheros AR928X wireless network adapter) [29] and hostapd daemon [30] is used for Wi-Fi AP implementation and management.

In order to evaluate the throughput gains achieved due to network-assisted multihoming compared with MPTCP two scenarios are deployed. An overview of the baseline setup for both the cases of MPTCP-based multihoming and MF-based multihoming is depicted in Fig.4. The underlying L2 connectivity is the same for both cases, except for the additional GNRS node in MF setup which provides the name-to-address mapping to the MF-enabled routers.

MPTCP v0.90 is installed on the end points (sender and receiver) [11] and for traffic generation in the MPTCP-based scenario, *iperf* has been used to measure the downlink application level performance. In order to measure the downlink throughput for MF-based scenario, *mfperf* is used, which is a modified version of iperf adopting MF API calls [31] to transmit and receive data. We ran 20-second data transmission, 10 times for each of the scenarios. For the sake of comparison, the same set of experiments are conducted for single LTE and single WiFi link connectivity, using iperf. The result of the average throughput at the client is shown in Fig.5.

As can be seen from Fig.5, single LTE and single WiFi have comparable throughput. However, performing ping test on each of the technologies revealed high disparity in the latency on each link (WiFi has latency of average 2 ms, whereas LTE's latency averages 40 ms). The higher latency in LTE is a result of the L2TP tunnel and the LTE's core network delay. The default scheduler in MPTCP is based on RTT of each path; it will firstly fill up the congestion window on the subflow with lowest RTT, and sequentially transmit on next-higher RTT. Previous work has shown that the latency disparity on links causes under-utilization of available resources due to reordering overheads [18]. This is also backed up by our experimental evaluation, which shows MPTCP can barely perform better than the best of single link data transmission.

On the other hand, Mobilityfirst is capable of better utilization of available resources due to features such as storage aware routers capable of temporarily caching in-transit chunks, along with the bifurcation router assigning chunks to each path based on their quality. This is validated through our experimentation which shows multihoming will achieve higher throughput gain compared with MPTCP.

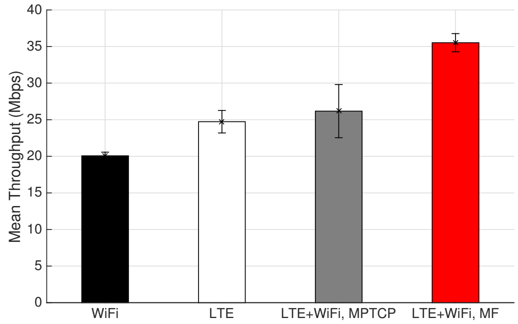

Fig. 5. Achieved downlink throughput at the client for 4 different scenarios: Single WiFi, Single LTE, MPTCP (LTE+WiFi), Moblityfirst(LTE+WiFi)

VI. FUTURE WORK

The MF evaluation described above using a small scale setup yields results in a localized environment but to truly understand the at-scale effect of the protocols we will need to deploy the baseline setup across geographically dispersed compute, network and storage resources to characterize different clients and communication paths. We plan to extend this baseline setup on the GENI wide area testbed with a last hop mobile edge. Each GENI deployment consists of a local cloud compute and storage cluster called a GENI Rack, which runs the EPC and uses high speed fiber connectivity to the LTE eNodeB [32]. GENI provides UE's in the form of USB dongles and Android handsets. These deployments are all interconnected over a research network backbone, Internet2 (see Figure 6). GENI also provides WiFi AP's that will allow us to evaluate multihoming (simultaneous connection to both WiFi and LTE) scenarios.

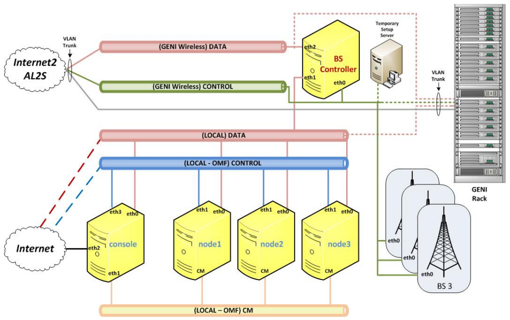

Fig. 6. GENI Mobile Edge Setup

VII. CONCLUSION

In this paper we present emerging mobility services in the context of 5G architecture, which would be necessary for 5G's promise to provide a significantly enhanced mobile user experience with Gbps wireless bit-rates along with low latency and improved reliability. Integrating these services enabled by Future Internet Architecture projects with 5G technologies and standards is an important step towards real-world deployment of these services. In this work a device multihoming scenario has been implemented using WiFi and LTE technologies in ORBIT testbed. The performance of mutihoming solution in Mobilityfirst FIA has been evaluated in comparison with current TCP/IP based solutions like multipath tcp. Through baseline throughput test experiments, the gains achieved through MF-based network-assisted multihoming have been illustrated.

REFERENCES

[1] M. Berman et al., "Geni: A federated testbed for innovative network experiments," *Computer Networks*, 2014.

[2] D. Raychaudhuri et al., "Overview of the orbit radio grid testbed for evaluation of next-generation wireless network protocols," in *IEEE Wireless Communications and Networking Conference*, 2005.

[3] "Nsf fia project," http://www.nets-fia.net.

[4] G. Bhanage et al., "Virtual basestation: Architecture for an open shared wimax framework," in *Proceedings of the Second ACM SIGCOMM Workshop on Virtualized Infrastructure Systems and Architectures*, 2010, VISA '10.

[5] G. Bhanage et al., "Vnts: A virtual network traffic shaper for air time fairness in 802.16e systems," in *2010 IEEE International Conference on Communications*, 2010.

[6] F. Fund et al., "Geni wimax performance: Evaluation and comparison of two campus testbeds," in *2013 Second GENI Research and Educational Experiment Workshop*, 2013.

[7] F. Bronzino et al., "Experiences with testbed evaluation of the mobilityfirst future internet architecture," in *2015 European Conference on Networks and Communications (EuCNC)*, 2015.

[8] R. Izard et al., *An OpenFlow Testbed for the Evaluation of Vertical Handover Decision Algorithms in Heterogeneous Wireless Networks*, 2014.

[9] Ch. Tsao et al., "On effectively exploiting multiple wireless interfaces in mobile hosts," in *Proceedings of the 5th International Conference on Emerging Networking Experiments and Technologies*, 2009, CoNEXT '09.

[10] K. K. Yap et al., "Making use of all the networks around us: A case study in android," in *Proceedings of the 2012 ACM SIGCOMM Workshop on Cellular Networks: Operations, Challenges, and Future Design*, 2012, CellNet '12.

[11] "Multipath tcp - linux kernel implementation," https://multipath-tcp.org.

[12] C. Raiciu et al., "How hard can it be? designing and implementing a deployable multipath tcp," in *Presented as part of the 9th USENIX Symposium on Networked Systems Design and Implementation (NSDI 12)*, 2012.

[13] Sh. Deng et al., "Wifi, lte, or both?: Measuring multi-homed wireless internet performance," in *Proceedings of the 2014 Conference on Internet Measurement Conference*, 2014.

[14] Y. Chen, "A measurement-based study of multipath tcp performance over wireless networks," in *Proceedings of the 2013 Conference on Internet Measurement Conference*, 2013.

[15] B. Sonkoly et al., "Sdn based testbeds for evaluating and promoting multipath tcp," in *2014 IEEE International Conference on Communications (ICC)*, 2014.

[16] A. Venkataramani et al., "Mobilityfirst: A mobility-centric and trustworthy internet architecture," *SIGCOMM Comput. Commun. Rev.*, 2014.

[17] P. Karimi et al., "Achieving high-performance cellular data Services with multi-network access," in *Proceedings of IEEE Globecom 2016 (to appear)*. IEEE, 2016.

[18] S. Mukherjee et al., "Network-assisted multihoming for emerging heterogeneous wireless access scenarios," in *IEEE PIMRC*, 2014.

[19] Sh. Mukherjee et al., "Achieving scalable push multicast services using global name resolution," in *Proceedings of IEEE ICC (to appear)*, 2017.

[20] M. Li et al., "Block-switched networks: A new paradigm for wireless transport," in *Proceedings of the 6th USENIX Symposium on Networked Systems Design and Implementation*, 2009, NSDI.

[21] "Mobilityfirst software release prototype documentation," http://mobilityfirst.orbit-lab.org/wiki/Proto.

[22] T. Vu et al., "Dmap: a shared hosting scheme for dynamic identifier to locator mappings in the global internet," in *Distributed Computing Systems (ICDCS), 2012*. IEEE, 2012.

[23] R. Morris et al., "The click modular router," in *ACM Transactions on Computer Systems*. Citeseer, 2000.

[24] S. Nelson et al., "GSTAR: generalized storage-aware routing for mobilityfirst in the future mobile internet," in *Proc. of MobiArch*. ACM, 2011.

[25] F. Bronzino et al., "Network service abstractions for a mobility-centric future internet architecture," in *Proceedings of the eighth ACM international workshop on Mobility in the evolving internet architecture*. ACM, 2013.

[26] Th. Rakotoarivelo et al., "Omf: a control and management framework for networking testbeds," *ACM SIGOPS Operating Systems Review*, 2010.

[27] M. Singh et al., "Orbit measurements framework and library (oml): motivations, implementation and features," in *Tridentcom 2005*. IEEE, 2005.

[28] "Openairinterface," http://www.openairinterface.org/.

[29] "Linuxwireless ath9k," https://wireless.wiki.kernel.org/en/users/drivers/ath9k.

[30] "hostapd: Ieee 802.11 ap, ieee 802.1x/wpa/wpa2/eap/radius authenticator," http://w1.fi/hostapd/.

[31] F. Bronzino et al., "Network service abstractions for a mobility-centric future internet architecture," in *Proceedings of the Eighth ACM International Workshop on Mobility in the Evolving Internet Architecture*, 2013, MobiArch.

[32] A. Gosain et al., "Enabling campus edge computing using geni racks and mobile resources," in *2016 IEEE/ACM Symposium on Edge Computing (SEC)*, 2016.