An OpenFlow Testbed for the Evaluation of Vertical Handover Decision Algorithms in Heterogeneous Wireless Networks

Ryan Izard, Adam Hodges, Jianwei Liu, Jim Martin, Kuang-Ching Wang, Ke Xu

Clemson University, Clemson SC 29634, USA
{rizard,hodges8,jianwel,jmarty,kwang,kxu}@clemson.edu

Abstract. This paper details a framework that leverages Software Defined Networking (SDN) features to provide a testbed for evaluating handovers for IPv4 heterogeneous wireless networks. The framework is intended to be an extension to the Global Environment for Network Innovations (GENI) testbed, but the essence of the framework can be applied on any OpenFlow (OF) enabled network. Our goal is to enable researchers to evaluate vertical handover decision algorithms using GENI resources, open source software, and low cost commodity hardware. The framework eliminates the triangle routing problem experienced by other previous IPv4-compatible IP mobility solutions. This paper provides an overview of the testbed framework, implementation details for our installation using GENI WiMAX resources, and a discussion of future work.

Key words: Heterogeneous Wireless Networks, OpenFlow, IP Mobility, Testbeds, Vertical Handovers

1 Introduction

Wireless networks and mobile devices have evolved to the point where access to multiple radio access technologies (RATs) are commonplace. The combination of these RATs form a heterogeneous wireless network. It is in the best interest of both users and network operators that the network resources be distributed fairly and efficiently amongst all users within a heterogeneous wireless network [3]. Vertical handovers, or handovers between different RATs, cause interruptions in connectivity for mobile users during the process of obtaining a different IP address. This results in the temporary loss of IP connectivity. Another issue presented by a vertical handover is that typical network applications are not designed to support the use of multiple network interfaces.

Mobile IPv4 (RFC 5944) [7] provides a mechanism by which a mobile device can retain the use of an IP address even after it has associated with a foreign network. Upon migration, the mobile device reports its new IP address to the home network, and a tunnel is formed between the mobile device and a home agent at the mobile device's home network. Egress traffic from the mobile device is routed normally, while ingress traffic is routed back to the home network and

through the tunnel to the mobile device as illustrated in Fig. 1. The use of a tunnel creates what is known as the Mobile IP triangle routing problem, which adds delay and consumes extra network resources. The introduction of IPv6 features to Mobile IP (MIPv6) [8] running on an IPv6 network can alleviate the triangle through the use of what is known as route optimization. A limitation of these MIPv6 schemes is that they require custom software on the client to enable mobility.
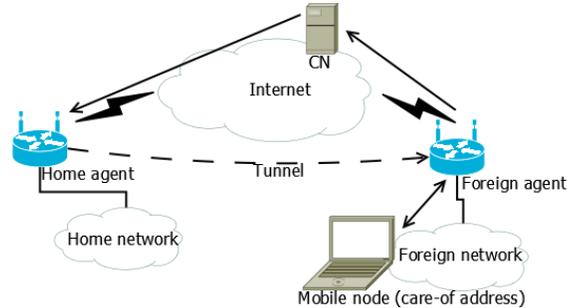

Fig. 1. Mobile IPv4 routing triangle problem

Many handover decision algorithms have been developed to alleviate the resource allocation problem presented by heterogeneous wireless networks. Most of these algorithms require the ability to make a handover decision based on the current conditions of all RATs. 802.21 [10] is a framework that was developed to support media independent handovers. The goal of this framework is to provide a standardized interface for every RAT that handover decision algorithms can utilize to create a global view of the available network states. An algorithm can build this global state by using 802.21 to query information, handle events, and issue commands. These algorithms can then use the global network state to intelligently trigger a handover. Testing vertical handover decision algorithms is nontrivial due to the cost and complexity of heterogeneous wireless network deployments and usually relies on network simulations [13]. We present a testbed design for evaluating these handover decision algorithms using real world wireless network resources.

The Global Environment for Network Innovations (GENI) [1] is a NSF sponsored effort to create a large scale testbed for network experimentation. One feature of GENI is its use of OpenFlow (OF) [2], a specification implemented by switches or routers that allows the forwarding plane to be modified by a controller in software. Recently, GENI has made strides towards enabling experimentation over wireless resources. It is a priority for GENI to provide mobile hosts with seamless vertical handovers using IPv4-compatible methods within GENI's heterogeneous wireless networks. This feature is desirable to researchers testing network applications as well as vertical handover decision algorithms.

Our contribution to this effort is the design of an OF-based, IPv4-compatible vertical handover testbed for use with GENI wireless networks. In the spirit of GENI, researchers will be able to reserve and utilize the resources in our implementation of this testbed. However, due to the nature of wireless research, we anticipate that other researchers may need an implementation of the testbed on their own campus. Therefore, the majority of this paper describes the details of how we applied, designed, and implemented the framework to provide a versatile testbed. An understanding of the challenges and issues we encountered and resolved will provide insight and guidance to other researchers who plan on implementing this testbed on their campus.

The remainder of the paper is organized as follows. In Section 2, we discuss and contrast our solution to similar testbed efforts. Section 3 outlines the goals, requirements, limitations, and user model of our testbed. Section 4 describes the system design and implementation details. Section 5 contains a discussion of our future work to be completed on our testbed. Finally, we conclude the paper in Section 6.

2 Related Work

In an effort to bring heterogeneous wireless network testbed resources to GENI researchers, GENI has partnered with Open-Access Research Testbed for Next-Generation Wireless Networks (ORBIT) [9]. GENI WiMAX resources have been established at a number of universities across the US and are available for researchers to use. These resources are linked via a L2 tunnel to the GENI testbed. A subset of these campuses have received ORBIT nodes that can be accessed by experimenters through the ORBIT Management Framework (OMF). Clemson is one of the GENI WiMAX campuses with ORBIT nodes. These nodes are ideal for serving as mobile hosts within a vertical handover-enabled heterogeneous wireless network testbed.

The WiRover [5] project at Wisconsin-Madison is a system that utilizes multiple radios to increase the bandwidth and continuity of wireless network access for buses. WiRover uses pre-collected signal data along bus routes to allow their system to proactively make an intelligent handover decision. Although the WiRover project itself is not a testbed, the researchers' previous efforts in [6] include using the vehicular network as a testbed for a 3G-WiFi heterogeneous wireless network.

There are several existing heterogeneous wireless network testbeds that provide vertical handover capabilities. For example, [4] uses MIPv4 to achieve IP mobility, enabling the researchers to evaluate novel handover decision algorithms. Other testbeds, such as [11], do not have the requirement of IPv4 compatibility and can make use of MIPv6 and its route optimization features.

The proposal in [12] conceptualizes the client component of our testbed that we have detailed in section 4.2 of this paper. In this proposal, the researchers describe how such a client could participate in a Handover as a Service (HaaS) scheme. With HaaS, a central database makes the handover decision for a mobile

host based on the mobile host's current location and historical network information for that location. We hope that our vertical handover decision testbed will eventually facilitate the implementation of the system discussed by the researchers in this proposal.

3 System Framework

The system framework utilizes OF to achieve IP mobility and application-transparent handovers. It is designed to be easily deployed to universities that have OF-enabled campuses or that can provide a subnet or VLAN for handover experimentation. The framework was developed and implemented at Clemson as a part of the GENI WiMAX project. Fig. 2 provides a general overview of how our testbed is constructed and how it integrates with the large-scale GENI testbed. At the architectural level, the framework allows OF-enabled mobile devices to roam across any OF-enabled wireless network. It requires a root OF switch as the testbed ingress/egress, as well as OF switches at each edge network. The challenge in building a testbed based on the framework is the required integration with existing network infrastructure at the campus. The testbed allows a mobile device to roam from wireless Network X (as illustrated in Fig. 2) to Network Y, preserving end-to-end socket connections that the device has with other hosts located either on the campus network or in the Internet. With these minimal assumptions, the testbed can be as simple or as complicated as desired. The design of our testbed can be divided into two major components: the network-level and the client-level.
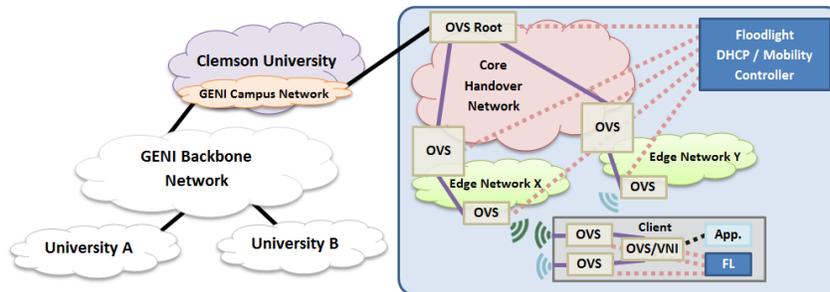

Fig. 2. Clemson testbed components and integration with GENI

The network-level component is required to manage and maintain client IP addresses, as well as the routing of client packets within the the testbed network. Both of these tasks are performed with a Floodlight (FL) OF controller, which maintains a global IP address pool and handles migration events within the testbed. To maintain the global IP address pool, a custom DHCP server module is integrated into FL. The FL controller is designed to be extensible to support other use cases; for example, in a HaaS framework, the FL controller could also

make the handover decisions for the mobile devices in the network [12]. A key component of the network-level is an OF-enabled switch or Open vSwitch (OVS) located at the root, such that all IP mobility-enabled networks on the edge are descendants of this root. As descendants of the root switch, other OF switches or OVSs are deployed on the network-level in order to both route client packets and detect migration events. From a network operations point of view, benefits of this tree-like design include (1) a single point of integration with the campus infrastructure and (2) the requirement of no specialized hardware in the case where OVS is used in favor over physical OF switches.

The client-level component of the testbed exists entirely on-board the client and is responsible for both switching the active physical interface and maintaining all client sockets during such a handover event. To maintain active sockets, a default virtual network interface (VNI) is installed on the client. All applications send and receive packets through this VNI, and by nature of a virtual interface, it is not impacted by physical interface states. The client is also equipped with OVSs and its own FL OF controller. This controller is responsible for routing packets from the VNI's OVS, through the client-level OVS network, and then to the physical interface of choice as determined by a handover decision.

Fig. 3. Clemson's GENI WiMAX/WiFi deployment

Researchers will soon be able to access specially configured GENI ORBIT nodes, both stationary and vehicular, through the ORBIT Management Framework (OMF) [9]. These nodes will be pre-configured to enable researchers to test their applications and handover decision algorithms in the seamless handover environment we have designed. Each ORBIT mobile node will be within overlapping coverage of both WiFi and WiMAX networks. Fig. 3 depicts the wireless coverage areas available to GENI researchers on the Clemson main campus.

4 Testbed Implementation

On the network-level of the testbed, all WiFi APs are configured with Debian Linux 5.1.10, and all WiMAX gateways are configured with Debian Linux 6.0.7.

On the client-level, testbed components have been verified on both Debian and Ubuntu Linux. All Linux distributions are using kernel 2.6.32. The FL OF controllers used on both the client and the network levels are sourced from FL 0.90. Each controller has been extended with custom modules to enable the vertical handover solution. Also common to both the network and the client are several OVS 1.9.0 virtual network bridges (OVSBs). A high-level diagram of the network-level is shown in Fig. 4.
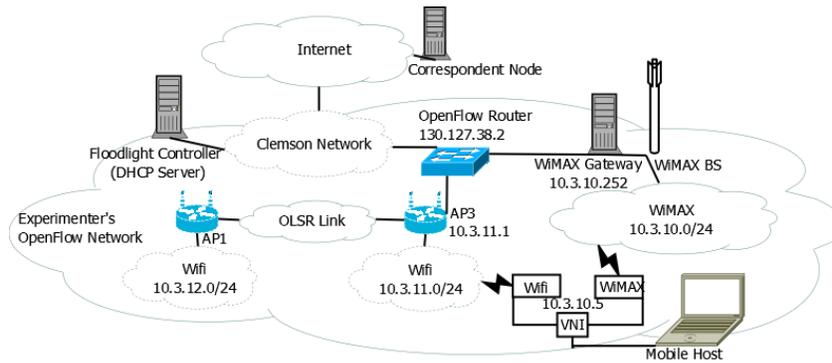
4.1 Network Component

Fig. 4. Clemson testbed network configuration

Within our testbed, the network component has the responsibility of maintaining the IP address pool for every mobility-enabled network. The network-level FL controller acts as a DHCP server, using DHCP requests as a trigger for migration. In the event of a migration, this Floodlight controller is also responsible for efficiently and quickly updating the client's location and thus the flow of its application packets.

The detection of a client connection and migration within the testbed is achieved through the use of OVSBs and OF flows (flows). These flows detect, encapsulate, and redirect client DHCP packets (on UDP destination port 67) to the network-level FL controller. This controller contains an integrated DHCP server module which, unlike traditional DHCP servers, associates an IP address lease with multiple MAC addresses. Each of these MAC addresses corresponds to a participating network interface (NIC) on the client. When processing a DHCP packet, the controller cross-references the MAC address of the DHCP packet with all available MAC address lists. Upon a successful match within a MAC address list, the controller assigns the client who initiated the request the corresponding IP. Then, upon a mobile host's initial connection or migration to a foreign network, flows are inserted in OVSBs starting at the testbed root and along every hop to the client's current location. These flows direct packets to and from the mobile client within the testbed. When a client migrates away

from this network, any existing flows associated with the client are removed and replaced with flows along the path from the root to the newly-migrated foreign network. The use of a root switch and tree hierarchy allows the network-level controller to avoid undesirable triangle-routing in the event of a migration.

The mobility testbed includes many OVSBs within the network. As discussed previously, the network-level OVSBs connect to the network-level FL controller / DHCP server. These OVSBs are used in the detection of client migrations and the routing of client packets into and out of the testbed. Specifically, the OVSBs on the testbed edge detect client migration by intercepting DHCP request packets, while the OVSBs in the core direct the flow of client packets from the testbed root to the client on the testbed edge.

Each network-level node with OVSBs also uses OVS patch ports (OVSPPs). To ensure proper routing of packets destined for an IP not routable by a foreign network, OVSPPs are used to connect the external and internal facing OVSBs installed on the gateways/APs. This allows independent subnets to operate within the testbed. The OVSPPs, combined with flows that utilize these OVSPPs, force client packets to bypass Linux routing on each hop, thus supporting cross-subnet compatibility upon migration from the home network.
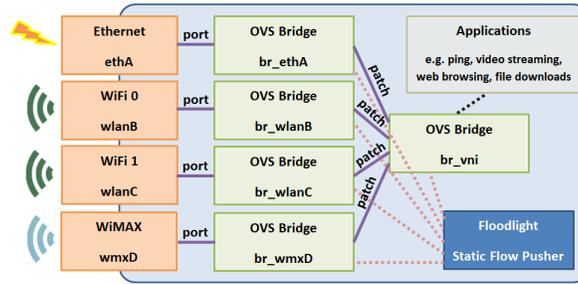
4.2 Client Component

Fig. 5. Client Open vSwitch interface configuration

Any mobile device should be able to connect to a network in our testbed and maintain an IP through a vertical or horizontal handover. However, if the handover is to be truly seamless to an application, there needs to be a persistent VNI for the application to use. The VNI abstracts the handover from the application and provides the application with an interface that is persistent for the duration the client is active. In addition to the VNI, the client should also be able to switch between interfaces in a manner that is simple and straightforward to the experimenter utilizing the testbed.

Similar to the network-level design, OVSBs are also utilized in the client to achieve a seamless handover. These client-level OVSBs are used in conjunction with a client-level FL OF controller and are installed for each mobility NIC on

the client, as shown in Fig. 5. The local FL controller inserts flows in each OVSB via the integrated Static Flow Pusher. These flows route application packets from the client VNI to the NIC of choice. When a decision is made to switch NICs, the client will issue a DHCP request egress the new interface, which will then trigger the aforementioned events in the network-level. As a result, these OVSBs with FL-inserted flows allow the client to seamlessly switch from one network to another. All client-level tasks are encapsulated in shell scripts to provide testbed experimenters with a simple and single command to execute a handover.

Each client-level OVSB also contain OVSPPs. To ensure proper routing of packets from the VNI to the NIC of choice, OVSPPs are used to connect the VNI OVSB with the OVSB of each participating NIC installed on the client. (as seen in Fig. 5). The OVSPPs, combined with flows that utilize these OVSPPs, serve to link the VNI to each NIC's OVSB. These flows define the route (and thus the NIC) used by application packets.

The use of a VNI introduces a problem when associating with networks and routing packets to the client from the network-level. The MAC address of the VNI must be the same as that of the NIC, otherwise WiFi APs and other access mediums will not accept packets from or know how to route packets back to the client's VNI at the link layer. It is not reasonable to require the modification or "spoofing" of each NIC's MAC address to that of the VNI. The client-level solution to this problem is to perform MAC-rewrite within the client OVSBs. When an application generates packets, they are routed out of the client via flows on each OVSB. These flows contain actions to rewrite the source MAC address of all egress packets from that of the VNI to that of the NIC. The flows also contain actions to rewrite the destination MAC address of all ingress packets from that of the NIC to that of the VNI. This rewrite process allows the client to send and receive packets from its VNI with any associated network on the link layer. Due to a limitation of OF, ARP packets cannot be rewritten with flows; they must be processed instead by a controller. Thus, the client-level FL controller contains a custom module to rewrite all ARP packet MAC addresses within the controller itself. Although out-of-band processing of packets is inefficient as compared to in-band, ARP packets are not frequent, so an occasional rewrite within the controller is a compromise made in our client-level implementation.

4.3 Client-Network Signaling

The sequence of events that takes place when a client migrates to a foreign network is shown in Fig 6. To summarize the interaction between the two components, the events are as follows:

1. Client establishes L2 connection with the network and issues DHCP discover
2. Floodlight intercepts packet, allocates IP and responds with an offer
3. Client responds to offer by sending DHCP request
4. Floodlight intercepts request, triggers migration event and DHCP acks
5. Client receives DHCP ack, establishing L3 connectivity. Meanwhile, the network-level FL controller inserts flows at the root and gateway OVSBs and removes any existing flows belonging to the client.

After this process completes, the client will have established full network connectivity on a foreign network through the root node.
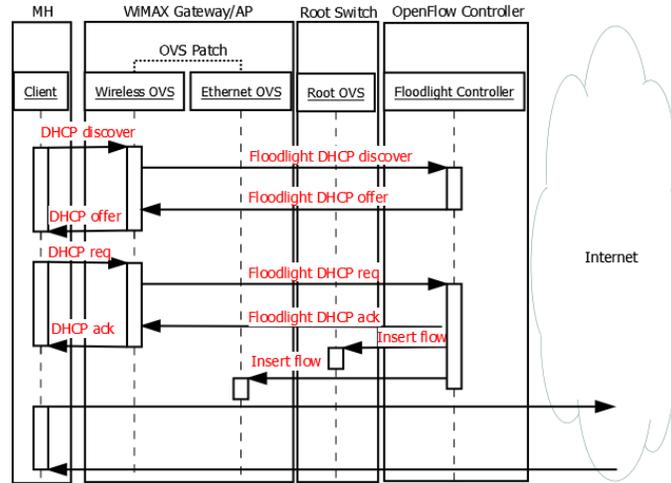

Fig. 6. Client Migration Sequence Diagram

5 Future Work

We plan on integrating our vertical handover solution more tightly with GENI wireless efforts. This includes helping other campuses install our testing framework as well as using OMF to manage our mobile nodes. The GENI wireless community has taken steps towards using Android handheld devices as their test devices for mobility scenarios. To accommodate this, we will be investigating installing Open vSwitch on GENI Android devices and deploying our client component as an Android application. This will enable researchers to test their handover decision algorithms on Android mobile devices in a real-world heterogeneous wireless network.

6 Conclusion

This paper details a framework that leverages SDN features to provide a testbed for evaluating handovers for IPv4 heterogeneous wireless networks. Based on the framework, we have designed and implemented a testbed that we (as well as other researchers) can use to explore ideas related to vertical handovers. The testbed was implemented as a part of the GENI project which identified the need for an easily deployable method of achieving seamless vertical handovers in an IPv4 OF-enabled heterogeneous wireless environment. The testbed design

presented in this paper meets these requirements and, as a proof-of-concept, has been implemented across wireless resources at the Clemson campus. The result is an IP mobility solution that is achieved solely through the use of OF features.

References

1. GENI (Jan 2014), `http://www.geni.net/`
2. OpenFlow (Jan 2014), `https://www.opennetworking.org/`
3. Amin, R.: Towards Viable Large Scale Heterogeneous Wireless Networks. Ph.D. thesis, Clemson University (2013)
4. Angoma, B., Erradi, M., Benkaouz, Y., Berqia, A., Akalay, M.C.: A vertical handoff implementation in a real testbed. Mobile Computing 1(1) (2012)
5. Hare, J., Hartung, L., Banerjee, S.: Beyond deployments and testbeds: Experiences with public usage on vehicular wifi hotspots. In: Proceedings of the 10th International Conference on Mobile Systems, Applications, and Services. pp. 393–406. MobiSys '12, ACM, New York, NY, USA (2012), `http://doi.acm.org/10.1145/2307636.2307673`
6. Ormont, J., Walker, J., Banerjee, S., Sridharan, A., Seshadri, M., Machiraju, S.: A city-wide vehicular infrastructure for wide-area wireless experimentation. In: Proceedings of the Third ACM International Workshop on Wireless Network Testbeds, Experimental Evaluation and Characterization. pp. 3–10. WiNTECH '08, ACM, New York, NY, USA (2008), `http://doi.acm.org/10.1145/1410077.1410080`
7. Perkins, C.: IP Mobility Support for IPv4, Revised. RFC 5944 (Proposed Standard) (Nov 2010), `http://www.ietf.org/rfc/rfc5944.txt`
8. Perkins, C., Johnson, D., Arkko, J.: Mobility Support in IPv6. RFC 6275 (Proposed Standard) (Jul 2011), `http://www.ietf.org/rfc/rfc6275.txt`
9. Rutgers, C., Princeton, Lucent Bell Labs, I.R., Thomson: Orbit lab. http://www.orbit-lab.org/ (Jan 2014)
10. Taniuchi, K., Ohba, Y., Fajardo, V., Das, S., Tauil, M., Cheng, Y.H., Dutta, A., Baker, D., Yajnik, M., Famolari, D.: Ieee 802.21: Media independent handover: Features, applicability, and realization. Comm. Mag. 47(1), 112–120 (Jan 2009), `http://dx.doi.org/10.1109/MCOM.2009.4752687`
11. Uddin, M.M., Pathan, A.S., Haseeb, S., Ahmed, M.: A test-bed analysis for seamless mipv6 handover in heterogeneous environment. In: Consumer Electronics (ISCE), 2011 IEEE 15th International Symposium on. pp. 89–94. IEEE (2011)
12. Xu, K., Izard, R., Yang, F., Wang, K.C., Martin, J.: Cloud-based handoff as a service for heterogeneous vehicular networks with openflow. In: Research and Educational Experiment Workshop (GREE), 2013 Second GENI. pp. 45–49 (2013)
13. Yan, X., ekerciolu, Y.A., Narayanan, S.: A survey of vertical handover decision algorithms in fourth generation heterogeneous wireless networks. Computer Networks 54(11), 1848 – 1863 (2010), `http://www.sciencedirect.com/science/article/pii/S1389128610000502`