Developing a Software Defined Networking Curriculum through Industry Partnerships

Carolyn J. Sher-DeCusatis and Casimer DeCusatis

***Abstract*—Software Defined Networking (SDN) is an emerging technology which radically improves cloud computing and other types of data networking. We discuss a new SDN undergraduate education program, developed in collaboration with industry partnerships. Student labs using resources such as GENI, NetFPGA, and the New York State Cloud Computing Center will be presented. We also outline SDN student projects including firewalls, load balancers, and redundant failover systems.**

***Index Terms*—Data Communications, Software Defined Networking, SDN, Network Function Virtualization, NFV, OpenFlow, Computer Engineering, Cloud Data Centers**

I. INTRODUCTION

The field of information technology (IT), particularly the design and management of data communication networks, is currently undergoing perhaps the most radical transformation in its long history. Instead of building and maintaining their own private data centers, many companies are becoming increasingly reliant on cloud computing (i.e. running their applications on servers and storage resources leased from a cloud service provider such as Amazon, Microsoft, or SoftLayer). According to industry analysts, today 20% of the Fortune 1000 use cloud computing for some applications, and half of them will be storing customer-sensitive data in the public cloud by 2016 [1]. A renewed emphasis on data networks for cloud computing is also disrupting conventional telecommunication companies and managed service providers (MSPs), who are rapidly trying to redesign their networks in order to remain competitive [2]. Applications such as social media, mobile access, video streaming, and big data analytics are driving this transformation at an accelerating pace. It has become challenging for traditional college curriculums to keep pace with the rapid developments in this area, and help insure their students will graduate with skills relevant to the modern IT job market.

This new, cost effective model for IT relies more strongly than ever before on dynamic, high performance networks within and between cloud data centers. A new approach to network administration has emerged known as software-defined networking (SDN) [3]. Although the term SDN was first introduced in 2009, there has always been some level of software control within data networks, so the exact meaning and impact of SDN has been the subject of much discussion and confusion in the industry. Beneath the hype, however, SDN offers significant technical advantages over conventional data networking.

SDN commonly refers to a set of technologies which separate the network control and management functions from the data transmission functions. Rather than requiring each switch to make hop-by-hop decisions on traffic routing, a centralized network controller provides end-to-end management of the data network. Further, SDN includes the abstraction or virtualization of network addresses and other functions. Just as server abstractions such as an operating system and programming language have superseded the requirement for all computers to be programmed in low level assembly language, SDN provides application programming interfaces (APIs) which automate many aspects of network management that are currently performed by low level, manual commands.

This is quite different from the approach used by traditional network switches and routers, which employ a distributed architecture (each switch only understands how to route data packets to the next hop in the network). Since there is no single point in the network which can enforce service levels, redundant connections, or network security, IT departments rely on network architects and administrators to approximately translate their business requirements into a series of provisioning commands which are manually entered into each switch on the network. Preparing for network administrative jobs in this environment traditionally involves complex, vendor-specific practitioner certification exams, which rely on memorizing switch configuration commands and learning how to implement hop-by-hop distributed networks. These skills are not well suited to an SDN environment featuring centralized, programmable network management interfaces.

Since the conventional network administration process is time consuming and expensive, most existing networks are statically configured. Changing these configurations is a

Manuscript received February 7, 2014. This work was supported in part by equipment donations from Columbia University and Stanford University, and software donations by Xilinx Corporation.

C.J. Sher-DeCusatis is an assistant professor of Computer Engineering Technology at New York City College of Technology, Brooklyn, NY 11201 USA (csher-decusatis@citytech.cuny.edu).

C. DeCusatis is a Distinguished Engineer with IBM Corporation, Poughkeepsie, NY 12601 USA. (decusat@us.ibm.com). He is also an assistant professor in the School of Computing Science and Mathematics at Marist College, Poughkeepsie, NY 12603 USA, and an adjunct associate professor of Computer Engineering Technology at New York City College of Technology, Brooklyn, NY 11201 USA.

978-1-4799-5233-5/14/$31.00 ©2014 IEEE

manually intensive process requiring days or weeks. By contrast, it is now possible to create several virtual machines (VMs) on one physical server in a matter of minutes [4]. SDN enables network functions to be dynamically provisioned as rapidly as virtual servers, enabling faster deployment of new applications, more efficient use of existing resources, and lower total operating cost. There is strong interest in applying the already successful principles of server virtualization to the network, and employers are highly motivated to find recent graduates with appropriate SDN skills.

By creating virtual networks running on top of existing physical switches, SDN makes it easier to replace physical devices (such as firewalls and load balancers) with their virtual equivalents, a process called Network Function Virtualization (NFV) [5]. By using NFV, it is possible to more easily implement "service chains", or traffic flows that link together different virtual appliances (for example, using SDN to route all traffic from a certain IP address through a virtual firewall before sending it along to a virtual load balancer). Both SDN and NFV are currently being used by many large corporations such as Google, Amazon, and Verizon [3, 6-7]. This has created a strong demand for network architects and administrators familiar with the theory and practice of SDN/NFV networking.

Conventional approaches to network education are not well suited to preparing students for this new environment. Historically, programming has been de-emphasized as a required skill for networking students [8]. It was felt that IT administrators didn't want or need to learn programming, which was more appropriate for traditional computer science degree programs. With the introduction of SDN/NFV, this approach to network administration education needs to change. For example, to provide high availability the network administrator no longer needs to manually provision two or more redundant physical switches. Instead, using APIs in an SDN network, it is possible to write a script instructing the network controller to dynamically fail over traffic to a path with lower service priority, and automatically switch back once the fault had been corrected. The ability to program network infrastructure APIs is rapidly emerging as a key differentiating skill for network architects and administrators, and will soon become a requirement for most employers. Successful graduates in this field require an understanding of SDN and NFV concepts, and how to apply these concepts to build and maintain automated, application aware networks.

There have been several reports on the need to reform engineering and computer science education [9], as well as reports on the transformative power of early curriculum redesign efforts in this field [10]. As part of this transformation, the gap between teaching methods and practitioner's skills can be addressed, at least in part, through industry and academic partnerships [11], which have been shown to help foster interdisciplinary education. We have found that this approach can also be beneficial to a practitioner's education program for cloud computing and SDN. While a number of leading academic institutions have recently begun to offer courses on SDN [12], it has so far been challenging to prepare students for IT careers in this rapidly evolving field, or to integrate these offerings into a more traditional undergraduate engineering curriculum.

Students in this field benefit from more practical, hands-on experience in creating SDN/NFV solutions, such as using virtual firewalls for network security instead of physical devices. Until recently, it was not cost effective to provide students with access to real world examples of IT infrastructure. New partnerships between industry and academia, as well as consortia between public and private universities, are proving to be an effective way to address this gap.

In this paper, we describe an approach to creating an undergraduate SDN/NFV curriculum being piloted in New York State, as part of an effort to encourage economic growth related to cloud computing and data analytics. We discuss implementation of this approach as a systemic way to increase the representation and advancement of underpriviledged and under-represented groups, such as the large, economically disadvantaged, nontraditional student population in the City University of New York. This approach includes not only technical skills but also promotes critical thinking, systems analysis, and interpersonal skills. We present a relatively low cost approach to SDN/NFV education, leveraging the nationwide GENI network and hands-on labs using the NetFPGA network interface cards. Finally, we discuss how this program benefits from a close alliance between CUNY, Columbia University, the State University of New York (SUNY), and Marist College (home of the New York State Center for Cloud Computing and Analytics), in addition to an ecosystem of active corporate sponsors and research affiliates.

II. SDN Concepts and Skills

A. Educational Goals

New York City College of Technology ("City Tech") is the designated senior college of technology within the 24-unit City University of New York (CUNY), the largest urban public university system in the nation. The college plays an important role nationally in the education of future scientists, engineers, technologists and mathematicians. The National Science Foundation ranks City Tech third nationally in the number of associate-level science and engineering degrees awarded to black students, 23rd in degrees awarded to male students, and 48th in degrees awarded to women. Baccalaureate programs are experiencing rapid enrollment growth; at present, 33.7% of all students are enrolled in baccalaureate programs.

Fall Semester 2012 student enrollment was 16,208, of whom 35% attended part-time. Approximately 31.5% of students self-identified as Black (non-Hispanic), 33.8% as Hispanic, 20 % as Asian/Pacific Islander, 11.3% as White, 0.6% as Native American, and 2.8% as Other. Sixty-one percent (61%) reported a household income of less than $30,000. Eighty percent (80%) of incoming first-year students and 65% of returning students received need-based financial aid. Sixty-seven percent (67%) are the first in their families to attend college. The student body reported 138 different countries of origin; countries of origin of faculty also span the globe. Nineteen percent (19%) of students reported working 20 or more hours per week. Students enter with widely disparate levels of academic preparation, professional goals, and personal circumstances. As an open access institution, City Tech's historic mission has been to offer opportunities for educational advancement to students regardless of financial circumstances or prior academic achievement. The college is a federally designated Hispanic Serving Institution (HSI).

The primary goal for CUNY students concentrating on data communications is to provide the background necessary to enable them to become successful Network Administrators, Network Architects, and Network Device Designers, within the context of a broader knowledge of computer engineering. We are particularly interested in local job opportunities in the nearby Wall Street financial district, where many employers are actively deploying SDN/NFV networks. There are a limited number of available hours in our curriculum that are not previously dedicated to other requirements, so it's important to prioritize key concepts and skills in networking education.

The fundamental concepts of SDN/NFV which our students should understand after successfully completing this course of study include the following:

- Similarities and differences between SDN and traditional network management, including distributed vs. centralized network management
- Abstraction and virtualization in SDN/NFV environments (for example, consequences of virtualizing layer 2/3 addresses and VLAN assignments)
- Programming SDN/NFV controllers based on open industry standards and APIs, and contributing to open source software projects on SDN/NFV
- Use cases in cloud computing, and how SDN/NFV affects performance
- SDN/NFV implementations using FPGAs, including Firewalls, Load Balancers, and Redundant Failover, as well as provisioning service chains (also known as function graphs)

B. Networking Classes

The computer engineering curriculum at City Tech allows students to earn a two year Associate of Applied Science (A.A.S.) degree in either Electro-Mechanical Technology, Electrical Engineering Technology, Telecommunications Engineering Technology, or Mechanical Engineering Technology. After completing two years of additional coursework, students can earn a B.Tech. degree in Computer Engineering Technology. These programs are ABET accredited.

SDN concepts are introduced as part of the fiber optics networking class. Skills such as critical thinking, troubleshooting, and system analysis are incorporated as part of the courses in Logic and Problem Solving as well as the Data Networking and Cloud Computing courses. B.Tech. classes related to this program provide a background including basic data network design, cloud computing, Python programming (used with the SDN controller), and FPGA programming (used with the SDN labs). Some existing course titles can be a bit misleading; the details of our integrated curriculum approach are given in Table 1. The fundamental courses, Cloud Computing and Electro-Optic Technology, are described further in the next section.

Table 1 – City Tech course sequence leading to SDN content

Course Title	**Content**	**Linkage to SDN**
Microcomputer Systems (required)	FPGA programming in various languages	Enables students to work with SDN labs emulating switches with FPGAs
Microcomputer Interfacing (required)	Prep for CCNA network certification exam	Provides background on conventional network design & management
Logic & Problem Solving (required)	Python programming	Python scripts will be used to manage SDN open source controllers
Applied Software Technology (elective)	Java Programming	Java is used to manage SDN open source controllers
Cloud Computing (elective)	IaaS and building cloud orchestration (Chef, Puppet)	SDN course content runs under IaaS
Data Networking (required)	Prep for network certification exams like A++ Net Cert exam	Provides background in fundamental network design
Electro-Optic Technology (elective)	Fiber optic networking and SDN labs	Primary focus for SDN content

III. INSTRUCTIONAL MATERIALS

The required textbook for the Cloud Computing class is *The OpenStack Cloud Computing Cookbook* by Kevin Jackson [13]. This book provides detailed instructions on how to build cloud infrastructure, including high level cloud orchestration tools such as Chef and Puppet. Labs for this course provide students with experience in building cloud computing environments including Infrastructure-as-a-Service (IaaS) [14]. This prepares students for related work in the SDN labs, which run beneath IaaS.

Previously, the Electro-Optic Technology course used a basic text on fiber optics which did not contain any SDN/NFV material (*Fiber Optic Communications* by Joseph Palais [15]. We have developed a supplement to this course (based in part on resources such as [16]) which explain SDN/NFV and their associated benefits. Portions of this course, and related courses such as Data Networking, are taught exclusively online as part of CUNY's distance learning program (instructors present live lectures online weekly, which are recorded so that students can view the content at any time). Distance learning students are required to attend labs on campus at least once per week. While the labs in the first half of the course are intended to develop a basic understanding of fiber optic networking, labs in the second half of the course are devoted to SDN concepts. The SDN labs we have developed use two different approaches. First, students have the opportunity to work with commercially available SDN equipment through the Global Environment for Network Innovation (GENI). Second, students collaborate in small groups for hands-on SDN device implementations using programmable NetFPGA hardware. We will discuss each of these approaches in turn.

A. GENI Network

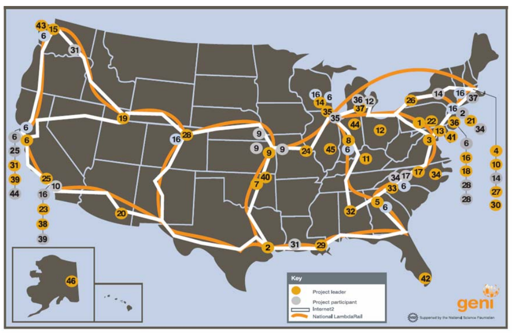

Figure 1. Overview of the GENI current generation "Spiral 3" nationwide network, which will support SDN using the OpenFlow protocols [17].

Established by the National Science Foundation and supported by a consortium of corporate sponsors, GENI [17] provides a virtual laboratory for networking and distributed systems research and education. GENI has created a nationwide network interconnecting Linux-based compute resources at major universities; a recent example of the GENI backbone is shown in figure 1. Academic membership is available at no cost, and provides access to this equipment and a supporting community who have developed training modules for many functions, including SDN. GENI also offers instructor training workshops, conferences, webinars, and online tutorials to facilitate incorporating their systems into an existing curriculum [17]. Participating researchers may download software into real or virtual GENI nodes, and experiment with at-scale, federated compute resources which would not be practical for a single university to maintain. As part of our CUNY lab program, students will complete exercises such as the GENI tutorials on firewalls and load balancing in an SDN environment. Students will also gain experience programming GENI nodes for different applications, and studying the interaction between servers, networks, and application software. Through the nationwide GENI network, City Tech is also able to collaborate with other academic institutions who share an interest in SDN education.

B. NetFPGA

Figure 2 – NetFPGA prototyping board [19]

Originally developed by Stanford University in 2007 as a means for rapid prototyping of computer devices [18], the NetFPGA is a line-rate, flexible, and open hardware and software platform for classroom experimentation and research. As the name implies, these devices use an FPGA-based approach to prototyping network switches and other devices. This platform is suitable for undergraduate education and research products related to SDN and data networking [19, 20]. As shown in figure 2, the 10 G NetFPGA card contains a Xilinx Vertex 5 FPGA chip (allowing programs written in high level programming languages to run on this card). The card also includes 4 x 10Gbps Ethernet Ports, as well as 27 MB QDRII-SRAM and 288 MB RLDRAM-II, and fits in a standard PCI Express port. There is also a NetFPGA opensource development community, which provides a repository for various device implementations.

CUNY presently has seven NetFPGA cards, donated by Stanford University, which are used in our SDN labs. Three to four students work together on each card. An SDN switch design for NetFPGA, implementing the OpenFlow protocol [19] has been developed by Stanford University and contributed to the open source repository. Students can download this and other design elements from the open source repository, or program their own applications using Xilinx ISE Design Suite software (donated by Xilinx Corporation). The NetFPGA cards have four connections for SFP+ optical or copper transceivers, which we have currently populated with short reach multimode optical transceivers (850 nm). This allows us to leverage the student's preparation in both FPGA programming and fiber optic technology. The NetFPGA cards can function as a small SDN switch, managed by an SDN open source controller. We also have an additional NetFPGA and computer for course development, on loan through our collaboration with Columbia University. We can give each team of students their own working copy of the Xilinx development environment.

C. Student Projects in Electro Optic Technology

The Spring 2014 class is the first to use SDN themed final projects in Electro Optic Technology. Each group of 3 to 4 students will complete the following assignments:

- Complete the GENI OpenFlow Tutorial

- Select a project theme: load balancing , firewall, or failover.
- Complete related tutorials in GENI, if available.
- Implement an example of the chosen theme using NetFPGA cards as programmable SDN switches.
- Produce a poster about this work, suitable for a technical conference, and present it to the class

We encourage students to create novel implementations or open source contributions to the NetFPGA libraries, and submit their work for presentation at professional conferences in our area which offer student poster sessions. This not only provides excellent experience for the students and promotes interpersonal communication skills, but also exposes them to potential employers in the region.

For example, consider the student project related to redundant failover. A traditional high availability network includes redundant physical switches and links, which are dedicated to the path between two application servers as shown in figure 3. High availability is achieved by configuring at least two dedicated, redundant physical links. When there is a failure in one path, the other continues to operate (there is no traffic switching between paths). To configure this environment, a network administrator would have to log into both switches, one at a time, to provision the necessary connections. This process is slightly different, depending on the switch manufacturer. Then the administrator needs to insure the physical links are statically configured just for this application. The administrator needs to manually re-configure traffic on both switches after a failure occurs and is subsequently corrected. This approach requires twice the hardware, footprint, and energy consumption as a non-redundant link. Since the redundant hardware sits unused most of the time, until there is a failure on the primary hardware, the switch is under-utilized and the installation cost for this service is high.

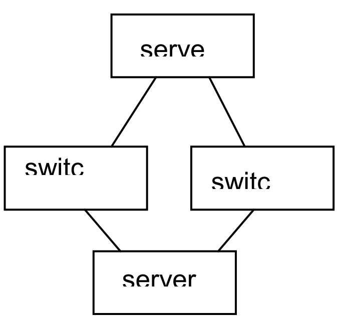

Figure 3 – Conventional high availability data network, statically configured

By contrast, an SDN high availability network can be quite different, as shown in figure 4. There may be multiple physical paths between two application servers; as long as there are at least two such paths, no physical reconfiguration of the network is required. High availability is achieved by dynamically re-configuring network paths in response to a failure. To configure the SDN environment, a network administrator would log into the centralized network controller, where she could write a programming script detailing conditions under which the link may fail over (this may include monitoring conditions outside the network, for example failures on the attached servers may initiate a network reconfiguration). This approach is the same for any manufacturer's switches in the network that support industry standard SDN. The failover is automated, and the controller can be programmed to revert to the original configuration when the failure is corrected. This approach may not require additional network hardware or physical reconfiguration; as long as there are underlying resources available, the SDN controller can reconfigure traffic under software control alone. Since the redundant path is switched on only when needed, and can carry other lower priority traffic at other times, the switches remain fully utilized and both the operating and installation cost of this service are reduced.

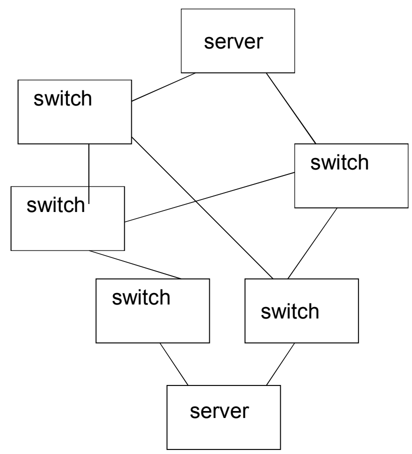

Figure 4 – SDN high availability data network, dynamically configured

For the SDN failover project, students learn about these differences and their implications. The students can configure both conventional and SDN failover scenarios on their NetFPGA cards (there is a demo of conventional failover on the NetFPGA community website). Students can appreciate the difference between manual and automated re-configuration, and see firsthand how writing a script for the SDN controller automates the failover process, saving time and insuring that the same steps are carried out every time this service is provisioned. By contrast, the conventional approach requires manual intervention every time there is a switch failure, and introduces the possibility of human error each time the switches need to be reconfigured.

IV. INDUSTRY-ACADEMIC COLLABORATION

Just as SDN/NFV is transforming the management of modern data networks, there is a need for exploring nontraditional approaches to SDN/NFV education and research. Specifically, SDN calls for a practitioner-oriented approach, and benefits from a closer interaction between educators and the IT administrators at their institutions. More meaningful collaboration between different branches of academia, or between academia and industry, would also benefit students in this field.

As part of an effort to promote economic development, in January 2013 the State of New York approved a $3M grant to establish the Center for Cloud Computing & Analytics at Marist College (a 4,500 student private college in Poughkeepsie, N.Y.). As part of this effort, Marist has established an SDN/NFV Innovation Lab, which serves as a test bed for next generation network virtualization. This test bed consists of three data centers interconnected using a 100 km optical fiber ring, with all networking equipment managed by a central SDN controller [21]. Network virtualization software is also used to create isolated, multi-tenant cloud computing environments.

Marist also exemplifies a unique collaboration between this academic education/research lab and their production IT systems, which take advantage of the research lab's experience with SDN technology to improve IT services to all students on campus. In keeping with their mission to promote the economic benefits of this technology across the state, Marist has formed academic partnerships with other public, private, and Ivy League schools, including CUNY, Columbia University, and the State University of New York (SUNY). Each of these institutions is conducting research on some aspect of SDN/NFV which serves an overall research agenda. For example, in addition to CUNY's work on SDN curriculum development, CUNY researchers are developing an analytic model to predict virtual machine migration times across an SDN network [21]. Columbia University students are working on approaches to perform real time traffic engineering for an SDN wide area network. Students at SUNY are developing a mobile management application for Android smart phones, in which context appropriate content from SDN network logs is delivered to a mobile app using messaging queuing middleware. This nontraditional, federated approach to technical education has yielded many benefits for the institutions involved, and provided students with a richer undergraduate experience.

Students at each of the participating schools can take advantage of the SDN test bed at Marist College to conduct undergraduate research projects or independent study, as well as developing a bridge to graduate studies. Remote access to the Marist test bed is being enabled for wireless devices, such as smart phones and tablet computers. Students are encouraged to develop open source SDN management application projects. Several undergraduate students working with Marist College have made meaningful contributions to open source SDN controllers, which are now used commercially [21]. For example, undergraduate students developed the software used in open source SDN controllers to detect when a physical switch port fails, and request graceful termination of all virtual traffic flows associated with that port (the so-called "port down reconciliation" problem). By making SDN accessible in this way, we can provide a much richer experience for undergraduate students with basic programming skills and an interest in data networking.

The SDN Innovation Lab is supported in part through corporate partners, which currently include Adva Optical Networking, BigSwitch, Ciena, IBM, NEC, and Plexxi. Other research partners include CERN (the European research organization based in Switzerland) through a partnership with the California Institute of Technology. These companies are not simply financial supporters of the lab; they are engaged in developing the research agenda, donating equipment for student use, and sponsoring student attendance at technical conferences. Corporate partners use the SDN Innovation Lab as a client showcase and briefing center, as well as an extension of their own development labs which can demonstrate first-of-a-kind SDN/NFV use cases. The collaboration between multiple industry sponsors and academic partners provides a force multiplier which increases the impact on a student's education, and is based on the National Science Foundation's Industry and University Cooperative Research Center (IU/CRC) model [22]. By training students with SDN principles that are of interest to these sponsors, this lab provides a very high placement rate for students after graduation. The success of this program has led to a collaborative effort among the various corporate and academic partners to develop a Massively Open Online Course (MOOC) for SDN, which is planned for initial availability in Spring semester 2014.

V. CONCLUSIONS AND FUTURE WORK

The industry-wide emphasis on cloud computing and data analytics has driven a renewed focus on networking technologies and the education process for network architects and administrators. SDN and NFV are becoming increasingly important, and skills in this area are in high demand. In an effort to bridge the gap between traditional network education, which does not include programming skills, and the emerging SDN/NFV environment, we have developed a novel curriculum at New York City College of Technology. This program is particularly well suited to engaging nontraditional and under-represented students because of its practical, hands-on focus and engagement with other academic and industry partners. Building on existing student skills in FPGA programming, fiber optics, and cloud computing, we have developed an SDN curriculum which can be deployed quickly at very low startup cost. We have made this technology accessible to a student population which includes a high percentage of under-represented students in the technical disciplines, enabling them to pursue opportunities as network administrators with leading financial companies and other employers. Lab experience using the GENI system and NetFPGA cards is important to delivering the maximum benefit to students in a limited time. Further, students have an opportunity to interact with their peers at other colleges in New York State, as well as with industry employers, through our collaboration with the New York State Cloud Computing

& Analytics Center at Marist College. Undergraduate students are capable of making meaningful contributions to research in this area, due to the emphasis on open source network controllers and industry standard management protocols. With this background, students are better prepared to create application aware, dynamically provisioned cloud data networks. In the future, we plan to produce more instructional materials, and add a capstone course solely devoted to Software Defined Networking (SDN) where students could work on full term research projects.

References

[1] IBM at OFC 2014: D. Freimuth, "Optical network requirements for cloud", Proc. OFC 2014, San Francisco, CA (to be published, March 2014)

[2] Communications Equipment Sector industry report 13-040, William Blair Group L.L.C.,, 27 pp., April 9, 2013 www.williamblair.com

[3] S. Das, D. Talayco, and R. Sherwood, "Software Defined Networking and OpenFlow", Chapter 17 in the *Handbook of Fiber Optic Data Communication*, C. DeCusatis, editor (4th edition), Elsevier/Academic Press (2013)

[4] J. Manville, "The power of a programmable cloud", OFC 2012 annual meeting, Anaheim, CA, paper OM2D.2 (March 18-22, 2013)

[5] Ibid L. Dunbar and D. Eastlake, draft IETF NV03 (DOVE) working group standard reference, Sept. 20, 2013 https://ietf.org/meeting/88/agenda/nvo3-drafts.pdf ; see also ETSI General Specification NFV 001 - 004 ver 1.1.1 (October 2013) http://www.etsi.org/technologies-clusters/technologies/nfv

[6] Sushant Jain, et. al., "B4:Experience with a Globally-Deployed Software Defined WAN", SIGCOMM 13, Aug 12-16, 2013, Hong Kong, China.

[7] D. McDyson, "Software defined networking opportunities for transport", IEEE Communications Magazine vol 51, no. 3, p. 28-31 (March 2013)

[8] I. Pepelnjak, *SDN, NFV and OpenFlow for Skeptics* (webinar and supplemental materials), Jan 23, 2014, http://cms.ipspace.net/bin/list?id=SDN101

[9] K. Jackson, *The OpenStack Cloud Computing Cookbook,* Packt Publishing, Birmingham, UK (Sept 2012).

[10] Kellog Foundation, "Unleasing new resources and entrepreneurship for the common good: a scan, synthesis, and scenario for action", January 1999

[11] "Extreme makeover college edition", Public Purpose, American Association of State Colleges and Universities, January 2010

[12] L.C. Wilcox and M.S. Wilcox, "A review and evaluation of engineering education in transition", IEEE Proc. 8th international conference on computer science and education (ICCSE), Sri Lanka (April 26-28, 2013)

[13] M. McNickle, "SDN Education", Tech Target magazine (posted July 2013, last accessed March 3, 2014), http://searchsdn.techtarget.com/feature/SDN-education-Universities-tackle-SDN-course-content

[14] Syllabus for City Tech course CET 4960, http://www.citytech.cuny.edu/academics/deptsites/cetech/docs/courses/CET4960.pdf

[15] J. C. Palais, *Fiber Optic Communications (5th Edition)*, Prentice Hall, Upper Saddle River, NJ (September 2004).

[16] IBM Redbook Solution Guide, *Connecting an IBM PureFlex System to the Network,* Published by IBM, Armonk, NY, January 30, 2013 (www.ibm.com/redbooks)

[17] J. Griffioen, F. Zongming, H. Nasir, W. Xiongqi, J. Reed, and C. Carpenter, "GENI enabled programming experiments for networking classes", Proc. Second IEEE GENI Research and Educational Experiments Workshop, Salt Lake City, UT, p. 111-118 (March 2013); see also GENI official website and wikis, http://www.geni.net/ ; http://groups.geni.net/geni/wiki/GENIExperimenter/Tutorials/OpenFlowOVS ; http://groups.geni.net/geni/wiki/GENIEducation/SampleAssignments/OpenFlowFirewallAssignment ; http://groups.geni.net/geni/wiki/GENIEducation/SampleAssignments/OpenFlowLoadBalancerAssignment

[18] J. Naous, G. Gibb, S. Bolouki, and N. McKeown, "NetFPGA: reusable router architecture for experimental research", In PRESTO '08: Proceedings of the ACM workshop on Programmable routers for extensible services of tomorrow, pages 1–7, New York, NY (2008); see also official NetFPGA website, http://netfpga.org

[19] J. Naous, D. Erickson, G. Adam Covington, G. Appenzeller, and N. McKeown, "Implementing an OpenFlow Switch on a NetFPGA Platform", ANCS '08, Nov 6-7, San Jose, California (2008)

[20] Daniel Wagner-Hall, "NetFPGA implementation of MOOSE", (undergraduate thesis), Homerton College, University of Cambridge, England (2010) http://www.cl.cam.ac.uk/~mas90/MOOSE/dwh-diss.pdf

[21] R. Cannistra, B. Carle, M. Johnson, J. Kapadia, Z. Meath, M. Miller, D. Young, C. DeCusatis, T. Bundy , G. Zussman, K. Bergman, A. Carranza, C. Sher-DeCusatis, A. Pletch, R. Ransom , "Enabling autonomic provisioning in SDN cloud networks with NFV service chaining", Proc. OFC Annual Meeting, San Francisco, CA (March 2014)

[22] National Science Foundation Industry/University Cooperative Research Center (last accessed March 3, 2014) http://www.nsf.gov/eng/iip/iucrc/