Virtual Public Networks

Arjuna Sathiaseelan[§] Charalampos Rotsos[§] Sriram C. S.[‡]
Dirk Trossen[§] Panagiotis Papadimitriou[†] Jon Crowcroft[§]

[§]Computer Laboratory, University of Cambridge, UK
{*arjuna.sathiaseelan, charalampos.rotsos, dirk.trossen, jon.crowcroft*}*@cl.cam.ac.uk*
[‡]Paxtera Solutions, USA
sriram@paxterrasolutions.com
[†]Institute of Communications Technology, Leibniz Universität Hannover, Germany
panagiotis.papadimitriou@ikt.uni-hannover.de

***Abstract*—Universal access to Internet is crucial. Several initiatives have recently emerged to enable wider access to the Internet. Public Access WiFi Service (PAWS) enables free Internet access to all and is based on Lowest Cost Denominator Networking (LCDNet) – a set of network techniques that enable users to share their home broadband network with the public. LCDNet takes advantage of the available unused capacity in home broadband networks and allows Less-than-Best Effort (LBE) access to these resources. LCDNet can enable third-party stakeholders, such as local governments, to setup, configure and operate home networks for public Internet access in cooperation with Internet Service Providers. Software-defined networking (SDN) creates new opportunities for the remote configuration and management of such networks at large scale.**

In this paper, we present Virtual Public Networks (VPuN), home networks created, deployed and managed through an evolutionary SDN control abstraction. This offers more flexibility to users and network operators, allowing them to share and control the network, while providing opportunities for new stakeholders to emerge as virtual network operators.

I. Introduction

Internet access in the recent years has become an important resource for the global population. The Human Rights Council considers Internet access as an important enabler of human expression and a potential human right [6]. However, the Internet is seriously challenged (infrastructural, socio-economical etc) to ensure universal access [8].

Lowest Cost Denominator Networking (LCDNet) [8] introduces a novel network paradigm for global Internet access, by utilizing unused network resources. LCDNet architects multi-layer resource pooling Internet technologies to support new low-cost access methods that could greatly reduce a network operator, as direct investment in local infrastructure to enable wider Internet access.

Amongst several initiatives for universal Internet access (e.g., [1], [2], [5], [11], [13]), Public Access WiFi Service (PAWS) [9] is based on LCDNet that makes use of the available unused capacity in home broadband networks and allows Less-than-Best Effort (LBE) [8] access (lower quality compared to the standard Internet service offered to paying users) to these resources. PAWS adopts an approach of community-wide participation, where broadband customers are able to donate controlled but free use of their high-speed broadband Internet to fellow citizens.

Large-scale deployment and management of such open networks will impose several challenges in terms of scalability, security, accountability and performance to both network operators and users. This will in turn increase operating expenditures for network operators to manage such networks. In [8] we argued that the stakeholder value chain should be extended in order to incentivize donated Internet access, by including more than the two traditional parties (i.e., ISP and consumer). The addition of third parties (e.g., local government or non-governmental organizations (NGO)) can in turn reduce the operational costs for ISPs. For users sharing their network, there will be major concerns with respect to security, performance, and network management. Home networks are already complex to setup and manage. This is a major obstacle for realizing our vision of wider deployment of a service like PAWS.

With the advent of software defined networking (SDN), there are more opportunities for network operators to deploy and manage in large scale such open public wireless networks. SDN has enabled open and programmable networks by isolating the control plane of the network and providing abstractions in it. In this paper, we use SDN for creating, deploying and managing such open public wireless networks which we define as *Virtual Public Networks (VPuN)*. Our architecture defines access point control abstractions that can be used by different stakeholders (users, network operators (NO) and third party virtual network operators (VNO)) to provide a third party VNO federated Internet access as well as the ability to dynamically control resources by both the NO and VNO.

The rest of the paper is structured as follows: In Section II, we discuss the requirements for the deployment of VPuN. Section III presents the VPuN architecture. In Section IV, we discuss the access network configuration required to deploy VPuN, Section V discusses the benefits of VPuN, while in Section VI we outline a security model for VPuN. Finally we conclude in Section VII.

II. Requirements

In this section, we briefly discuss the requirements for VPuN deployment and configuration. VPuN are envisaged to achieve the following high-level objectives:

- Expose a dynamic and user-friendly abstraction to stakeholders at various levels of the network to specify network resource requirements.
- When external factors compel it, VPuN should be automatically set up or reconfigured. This will be useful during emergency situations or natural disasters, where access points can automatically mesh with access points of other users or personal devices based on online social networking trends, or when network operators need to dynamically allocate capacity to meet evolving demands (e.g., flash crowds).

To fulfil these high-level objectives, VPuN must:

- Provide the ability for stakeholders to specify their requirements and the parameters that affect these requirements in a simple manner.
- Provide the ability to translate these requirements into control flows to be installed in various components of the network.
- Aggregate and curate data from authorized social media and news feeds and form the network itself to ascertain status of the environment.
- Using the information gathered, make intelligent decisions to automatically reconfigure the network.

In addition to these feature requirements, VPuN must fulfil the following design requirements so that it can be integrated seamlessly with any SDN stack.

- Provide backwards compatibility with existing network protocols.
- Provide ability to extend and support other SDN specifications that might evolve over and above OpenFlow.
- Provide transparent APIs and libraries that can be used to build external value-added applications (e.g., a reward point tracker for capacity sharing).

VPuN require the following components:

- SDN enabled home routers that can be configured using open APIs (e.g., OpenFlow [7]).
- Controllers (e.g., POX, NOX [4]) that allow users to access and modify the flow table of a home router.
- A sharing policy expression language (SPEL) that allows each home network user to specify, amongst others, the amount of bandwidth and the period over which his network will be shared with VPuN clients.

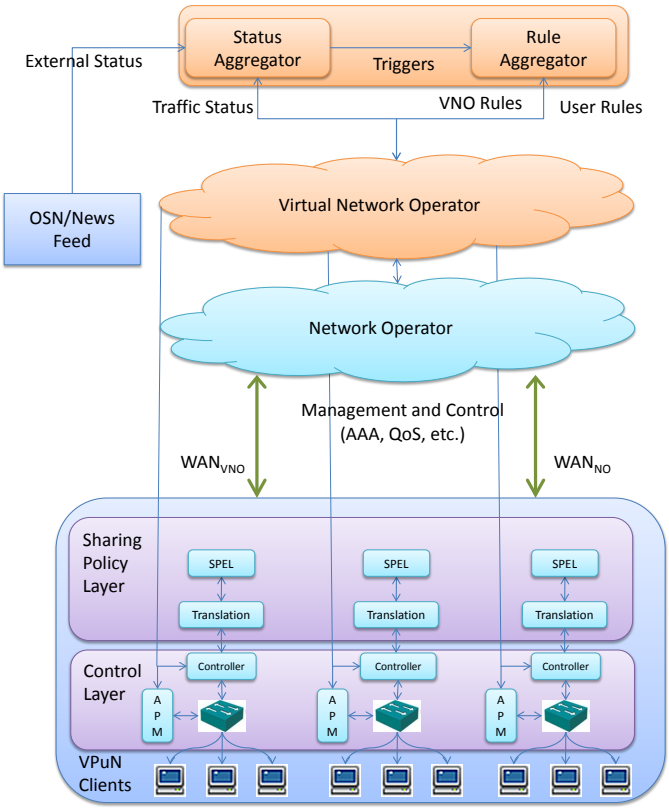

Fig. 1. VPuN architecture overview.

III. Architecture Overview

In this section, we present an overview of the VPuN architecture. As illustrated in Fig. 1, VPuN architecture is distributed across home networks, the network operator (NO) and the virtual network operator (VNO). According to the VPuN requirements, an SDN-enabled home router, a controller, a Access Point Manager (APM) and a system for the translation of sharing policy expressions are deployed in each home network. The APM is the system that contains the router control API described in Section III-D. In addition to the basic APIs described there, the APM will also provide the ability to create, modify and configure wireless SSIDs. As such, the APM and home router can be configured using a SDN controller (e.g., NOX [4]). An alternate approach would be to use a slicing layer (e.g., FlowVisor [10], [14]) for the home router allowing both the home network user and the VNO to control their respective slices while enforcing isolation between these slices.

The VNO also deploys the status and rule aggregator whose functionality is discussed below.

A. Sharing Policy Expression

We propose a sharing policy expression language (SPEL) for the specification of network sharing policies. As such, a

user can specify the amount of bandwidth that is willing to contribute and the period that he desires to share his network. Network sharing may be subject to other conditions and events which can be comprehensively expressed using SPEL. Besides the language specification, VPuN requires the translation of sharing policy expressions into network programming actions sent to the SDN controller.

SPEL employs an XML-based schema, as shown below:

$< rule\_id > 1 < /rule\_id >$
$< name > Share\ when\ I\ sleep < /name >$
$< condition >$
$time\_of\_day > 8.00 PM\ AND\ time\_of\_day < 5.00 AM$
$< /condition >$
$< action\ value = SHARE >$
$< data\_cap\ value = 1 >$
$units = GB$
$< data\_cap >$
$< rate\ value = 4 >$
$units = Mbps$
$< rate >$
$< /action >$
$< expire > 00:00:00\ 31\ Nov\ 2013\ IST < /expire >$
$< priority > 1 < /priority >$
$< watch\ value = "My\ event\ cancelled" >$
$< action > EXPIRE < /action >$
$< handle\ value = @AuthorizedHandle >$
$< type > Twitter < /type >$
$< /handle >$
$< /watch >$

As shown in the example above, a rule consists of the following clauses:

- A condition clause that tells what condition must be met for this rule to become active. This clause includes variables, such as $time\_of\_day$, $current\_location$, my_subnet, my_vlan_id, and comparison operators to check for greater than, lesser than, equals to and matches.
- An action clause that will state what should be done when the condition is met. Common actions comprise $SHARE, DENY$, for sharing capacity on the access point created by the access point manager or denying traffic. The sub clauses for an action will be $rate$ and $data\_cap$ followed by the associated units.
- A watch clause that tells the SPEL what expression to watch for in authorized handles. The handle is specified as a type (e.g., Twitter, Facebook, News) along with an id. The watch clause sets triggers in the status aggregator, as described in the next section, resulting in rules being installed automatically. An action sub clause is added to this watch clause stating what must be done when a certain watch pattern is found on the handle. The action sub clause supports an extra action other than those listed above which is $EXPIRE$. This action will cause the rule to expire immediately. In addition to social media handles, an additional handle type of $network\_info$ is also supported to listen for network events and information such as node down and capacity updates.
- Besides these, certain parameters like expiry timeout, priority to resolve conflicts, $ruleid, name$, etc., will also be supported.

A VNO aggregates all the rules installed by various users for the purpose of analyzing trends and managing the network. The knowledge of users' sharing policies is a prerequisite for the VNO to assign guest users to access points aiming to balance the load among open home networks.

It is not required that the SPEL translator and the SDN controller are collocated with the home router. Instead, they can reside on a server hosted by the NO or the VNO.

B. Status Aggregator

The status aggregator is a curated service that gathers information about the external world, such as natural disasters, crowd distribution, weather changes, via social media and other authorized news feeds, and the network, such as network load, traffic conditions and patterns, load distribution, via distributed controllers (for SDN-enabled network components) or SNMP (for traditional networks). The information that is scanned for is determined by filters applied by the rule aggregator (Fig. 2).

Data received from various state providers is combined into a unified JSON/XML format and provided to the VNO. This information can be monitored manually by the VNO so that corrective actions can be taken in situations that demand them. Alternatively, the VNO can also enable a rule aggregator which takes into account all rules in the system and takes corrective actions automatically.

Note that the status aggregator can be a distributed service that can split aggregation across individual status aggregator nodes and share status data, thereby enabling efficient use of resources.

C. Rule Aggregator

The rule aggregator is an intelligent decision making engine that can act as a proxy agent for an end user or VNO, make informed decisions and install rules in the system for them. The rule aggregator is just an extension of the SPEL with two additional functions: (i) a listener port that receives data from the status aggregator, and (ii) a listener port that consolidates rules from distributed client rule engines.

Based on the rules aggregated from VNOs and end users, the rule aggregator pushes filters to the status aggregator(s). The status aggregator in turn pushes triggers to the rule aggregator to install rules automatically as and when the system demands it. This allows the rule aggregator to scan the network and external world for relevant information and make informed decisions on behalf of both the VNOs and the end-users.

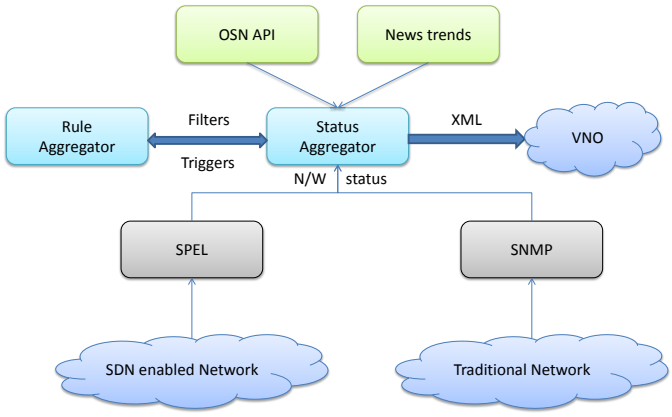

Fig. 2. Status aggregator flow.

Similar to the other components, the rule aggregator can be centralized or distributed as per deployment demands.

D. Router Control Abstraction

In our current VPuN design, we employ the home router as the primary mechanism to enforce the sharing policy of the network. This design approach simplifies system deployment, requiring solely a home router upgrade. VPuN home router comprises of a switching mechanism, exposing an OpenFlow control interface, a custom OpenFlow controller and a control interface for the local household network. The VPuN controller redefines the router control architecture and exposes a minimal JSON-RPC API, enabling third-party VNOs access and resource allocation policy configuration and usage logging. In addition, the switching mechanism must expose to the controller a configuration API for traffic shaping queues and Access Point (AP) management.

VPuN requires an isolated wireless network within each home network. In the data-link layer we enable virtualization through the multi-SSID functionality of modern wireless chips. The home router exposes two distinct wireless networks: one for the home network and the other for the VPuN (for e.g. PAWS). In addition, we separate data plane traffic through control plane virtualization in the network controller. We use the SPEL defined in Section III A for allowing home users to specify network sharing policies and transforming them into OpenFlow-based policies. Alternatively, the router can also use existing home network control interfaces to enable homeowners to express local network policy and transform them into appropriate OpenFlow-based policies [15]. Nonetheless, the user can exercise control only on traffic under the home network subnet, originating either by the home wireless interface or the broadband link interface.

For VPuN control, we modify the network control logic in order to develop a new abstraction that matches VPuN management requirements. We expose through the controller a JSON-RPC API to enable VNO policy expression. The API comprises of four functions, presented in Table I. $User\_On$ and $User\_Off$ functions are implemented by the VNO services and provide guest user connection and disconnection notification. In order to enable Internet connectivity for a guest user, the VNO must modify router policy, using the $User\_Config$ method. The method is used to specify guest user service accessibility (e.g., permitted domain name or IP addresses) and resource aggregate billing information on a per user basis, using the Get_User_Stats method. The method provides accounting and accountability information per guest user to the VNO.

TABLE I
ROUTER CONTROL API

Function	Description
$User\_On$ / $User\_Off$	Router → VNO: notify VNO on the connection/disconnection of an authenticated user
$User\_Config$	VNO → router: VNO configures resource parameters for the user
Get_User_Stats	VNO → router: VNO requests traffic usage information

The proposed control architecture increases the forwarding complexity of the home router. Nonetheless, the low traffic rates on the edges and the forwarding flexibility of software home routers does not affect the local network performance [17].

E. Third-Party Applications

While technically not a part of the architecture, all the components in the architecture will expose APIs that can be used to build external applications that can provide value-added services to end-users and VNOs. Examples of such applications are the following:

- A reward point management application that tracks information received from the rule aggregator and status aggregator to decide actual bandwidth shared by users and reward them redeemable points.
- A social media based access enabler that determines APs that belong to friends.
- An intelligent AP identifier that talks to local status and rule aggregators to determine which of the free APs will be available based on e.g., most bandwidth, highest duration (based on rules), VNO/NO preference.

While the architecture shows these applications on the top of the stack, these applications could reside anywhere from the VNO to a mobile device of an end user.

IV. ACCESS NETWORK CONFIGURATION

Enforcing resource control policies on the home network is not sufficient for end-to-end resource allocation. In the majority of current broadband networks, network bottlenecks affecting traffic prioritisation, occur in network points beyond the control of the end-user, within the backhaul of the ISP network. In order to provide accurate short-term end-to-end

resource allocation, we need to enforce resource allocation policies within the ISP network.

Our design provides an evolutionary deployment mechanism. In order to establish clean separation between the home user and VPuN traffic in the ISP network, the NO must support multi-addressing for each household. Specifically, each access point has a public IP (WAN_{NO}) per household and multiple private IP addresses WAN_{VNO} (depending on the number of VNOs), which are used to route traffic for VPuN clients (the VPuN wireless interface is bridged to the respective WAN_{VNO} interface). The NO through its routing policy aggregates incoming and outgoing traffic in a single point within the network to apply NAT translation, as well as, enforce per-VNO aggregate resource allocation and accounting. Since the VNO traffic constitutes an aggregated IP subnet, the NO can apply QoS polices to mark all VPuN traffic at a lower QoS (Less than Best Effort (LBE)) compared to paid user traffic (Best Effort (BE) or higher) or higher QoS depending on the established SLA between the NO and the VNO.

A SDN-enabled router can trivially support this policy mechanism. OpenFlow control provides primitive to translate source and destination IP addresses on a per-flow basis, thus tagging traffic on the end-nodes. In addition, the controller is responsible to enforce resource control for each user. The $User\_Config$ configuration API call contains rate limiting configuration parameters, which are translated to per-user queue setup. Using the $User\_Config$ API, the VNO can coordinate the assignment of VPuN clients to access points based on information available to rule aggregators and status aggregators, such as traffic loads. In case the VPuN client is disconnected, a connection with another proximate access point can be re-established, as shown in recent work [12].

V. Benefits of VPuN

We see four beneficiaries for our proposed architecture, namely the NO, a possible VNO, application providers, and the end user(s). As for the NO, we see three drivers for adopting our solution. Firstly, the introduced LBE service class allows for capitalizing unused bandwidth via the VNO towards customers who would otherwise be left out (while we can assume that a certain percentage of BE users would shift to the likely cheaper LBE class, we still assume that the overall LBE user base is larger than this shifting user base). Secondly, the proposed architecture allows for extending the overall customer base without new deployments beyond those made for the BE class i.e., the LBE class piggybacks on BE deployment by virtue of changing the control plane architecture. The extension of the user base, and therefore the additional income, can be used to offset the necessary capital investment into the proposed architecture. Thirdly, this expansion of end user reach comes likely at less cost in terms of energy efficiency. Although significant evaluation work is still outstanding for verifying this claim, we do expect that the re-usage of otherwise wasted bandwidth will come cheaper than merely deploying more bandwidth (which is again wasted at many times of the day). With that, the NO can aim at fulfilling possible regulatory requirements in terms of (a) user base and (b) energy caps for the deployed infrastructure through the introduction of the LBE class.

As a VNO, we see in particular local government agencies, charities or grassroot user communities implementing this role. Not only will this VNO provide these groups with the potential to reach end users who will otherwise be left out, and therefore address societal objectives, but it will also be doing so at lower price points than today's BE deployments. The latter might also be important for environmental objectives that these groups have, i.e., capping the necessary energy expenditure for such expansion of the Internet user base by piggybacking on existing deployments. Furthermore, we see room for so-called triple bottom line accounting approaches [18], where the economic objective of running of the VNO with an economical viability can be complemented by societal and environmental objectives. Although such interwining of societal and economic goals is generally doable at the level of the NO, its incorporated structure often emphasizes the economic dimension, despite notable efforts in increasing corporate social responsibility (in other words, we have yet to see an Internet deployment by a major ISP that is largely societal objective driven rather purely commercially motivated).

We expect that VPuN will be very useful during emergencies, such as natural disasters or terrorist activities, since home networks could be opened up to the public to communicate. The VNO has the choice on which access points can be opened up: the home user could have agreed to open up his entire network during such events. Using the external data feeds from OSN and news trends, the home networks can then be opened up either through manual VNO setup or automatically. VPuN could also be used to transfer government and industry sensor data without the need for building new network infrastructure thus reducing capital expenditures. New revenue models can be generated in such a way that sharers are provided with financial incentives.

Application providers are the third group of beneficiaries in our proposed solution. With our proposed solution to executing capacity dimensioning rules in the network, we can see space for application providers to insert service differentiation, such as upload services for media, that shift certain traffic from the BE class to the LBE class, utilising for instance delayed transmissions and additionally provided storage. Such service differentiation in turn can be offered to end users to deal with situations of temporary overload (e.g., in a large crowd situation) when trying to upload their precious media to, e.g., social networking sites. With separating the rule execution between core and access network, we can foresee such services being also core network facing, e.g., by limiting the overall usage of BE traffic at peak time without differentiation to specific users but an entire application.

Last but not least, end users are certainly a beneficiary of our proposed solution, not only through extending the reach of connectivity beyond the BE class users. We also see end users involved in the creation of the VNO role through necessary campaigning for establishing, e.g., local

user communities or dedicating local agency funds for such VNO role. In other words, end users can directly become engaged in the extension of the Internet by making themselves heard as potential new users and creating a case for VNO participation on both economic as well as social grounds. VPuN also removes the current concerns of users to share their home network infrastructure with the public by allowing them to have more finer grain control of their home network easing network management and at the same time providing strict traffic isolation, security and performance. This would enable wider adoption of a service like PAWS and hence more wider network coverage thus providing more opportunities for universal Internet access.

VI. Security Model

Open wireless infrastructures raise significant privacy concerns for users. We aim to provide two important privacy properties: strong user authentication and traffic isolation. Strong user authentication is required by the VNO to enable accurate traffic billing and user identification in cases of malicious behaviour. Traffic isolation is required to avoid traffic eavesdropping by collocated devices.

We propose the usage of the standardised WPA security mechanism to encrypt traffic between the device and the router within each VPuN wireless shared network. WPA provides pair-wise encrypted channels between the home router and each device, thus reducing the effectiveness of an eavesdropping atack. User authentication and pair-wise key generation is performed using the EAP-TLS mechanism [3]. EAP-TLS uses an X.509 public cryptography mechanism to establish a secure channel between the device and the AP and negotiate symmetric cryptography keys for the data plane traffic. Each VPuN network must provide a Certificate Authority (CA) (i.e., the VNO) and each AP point and device must use during authentication a certificate signed by the CA. Each VPuN certificate contains user information in the metadata, which is used to associate users with devices and network flows. Further, we advocate the integration of the Controller with the AP authentication mechanism. The controller receives a notification from the AP, every time a new authenticated device is connected to the network and propagate this information to the VNO.

Furthermore, in order to mitigate potential local network scanning attacks between devices within a VPuN, we reuse a series of network layer protocol intervention, proposed in [16], to exercise flow-level traffic control and device isolation for the VPuN network. For each device we use the DHCP protocol to partition the VPuN subnet in /30 prefixes and assign each device in an isolated subnet. Using these techniques, the network controller is able to suppress communication between devices belonging to different users, while broadcast discovery probes like ARP will never be forwarded to other devices.

Securing public networks faces a fundamental trade-off between performance, privacy and user-friendliness. We opt for a mechanism that enables security only in the first hop of the system. Our design is susceptible to privacy attacks, if the router is compromised. Similarly, we could establish a secure channel between the user device and the VNO and route traffic through the VNO network (for e.g. setup a Virtual Private Network (VPN)). We believe that this approach will centralise packet forwarding and incur significant network latency and jitter.

VII. Conclusions

With the advent of SDN, there are more opportunities for network operators to create, deploy and manage open home networks at large scale. In this paper, we presented Virtual Public Networks (VPuN), home networks created, deployed and managed through an evolutionary SDN control abstraction. We also discussed that the proposed control abstraction enables more flexibility for users and network operators to share and control the network and to allow new stakeholders to emerge as VNOs.

VIII. Acknowledgments

This work was partially supported by the EPSRC Grant EP/K012703/1. We would like to thank Dr. Richard Mortier and the anonymous reviewers for their feedback.

References

[1] Airjaldi, http://drupal.airjaldi.com.
[2] G. Bernardi, P. Buneman, and M. K. Marina, Tegola tiered mesh network testbed in rural Scotland, ACM MobiCom 2008 Workshop on Wireless Networks and Systems for Developing Regions (WiNS-DR'08), September 2008.
[3] P. Funk and S. Blake-Wilson, Extensible Authentication Protocol Tunneled Transport Layer Security Authenticated Protocol Version 0 (EAP-TTLSv0), RFC 5281, August 2018.
[4] N. Gude et al., NOX: Towards an operating system for networks, SIGCOM CCR, Vol. 38, No. 3, July 2008, pp. 105-110.
[5] S. Hasan et al., Enhancing rural connectivity with software defined networks, ACM DEV, January 2013.
[6] F. La Rue, Report of the Special Rapporteur on the promotion and protection of the right to freedom of opinion and expression, Human Rights Council, UN General Assembly, 16 May 2011, http://goo.gl/MDjS7
[7] N. McKeown et al., OpenFlow: Enabling Innovation in Campus Networks, ACM SIGCOMM CCR, Vol. 38, No. 2, pp. 69-74, April 2008.
[8] A. Sathiaseelan and J. Crowcroft, LCD-Net: lowest cost denominator networking, ACM SIGCOMM CCR Vol. 43, No. 2, pp. 52-57, April 2013.
[9] A. Sathiaseelan et al., Public Access WiFi Service (PAWS), Digital Economy All Hands Meeting, Aberdeen, October 2012.
[10] R. Sherwood et al., Can the Production Network Be the Testbed?, USENIX OSDI, Vancouver, October 2010.
[11] S. Surana et al., Beyond pilots: Keeping rural wireless networks alive, USENIX NSDI, April 2008.
[12] L. Suresh et al., Towards programmable enterprise WLANS with Odin, ACM SIGCOMM HotSDN, August 2012.
[13] Technology For All, http://tfa.rice.edu/.
[14] Y. Yiakoumis et al., Slicing Home Networks, ACM SIGCOMM HomeNets, Toronto, August 2011.
[15] M. Chetty et al., Refactoring network infrastructure to improve manageability: a case study of home networking. ACM SIGCOMM CCR, Vol. 42, No 3, pp. 54-61, July 2012.
[16] R. Mortier, et al, Control and understanding: Owning your home network, Communication Systems and Networks (COMSNETS), Jaunary 2012.
[17] C. Rotsos et al, Oflops: An open framework for openflow switch evaluation, Passive and Active Measurement, Springer Berlin Heidelberg, March 2012.
[18] Triple bottom line, $http://en.wikipedia.org/wiki/Triple\_bottom\_line$.