

# Flexible Traffic Management in Broadband Access Networks using Software Defined Networking

Julius Rückert<sup>\*1</sup>, Roberto Bifulco<sup>†1</sup>, Muhammad Rizwan-Ul-Haq<sup>\*†1</sup>, Hans-Joerg Kolbe<sup>†</sup>, and David Hausheer<sup>\*</sup>

<sup>\*</sup> Peer-to-Peer Systems Engineering, Technische Universität Darmstadt, Germany

Email: {rueckert|hausheer}@ps.tu-darmstadt.de, muhammad.rizwan-ul-haq@stud.tu-darmstadt.de

<sup>†</sup> NEC Europe Ltd, Heidelberg, Germany

Email: {roberto.bifulco|kolbe}@neclab.eu

**Abstract**—Over the years, the demand for high bandwidth services, such as live and on-demand video streaming, steadily increased. The adequate provisioning of such services is challenging and requires complex network management mechanisms to be implemented by Internet service providers (ISPs). In current broadband network architectures, the traffic of subscribers is tunneled through a single aggregation point, independent of the different service types it belongs to. While having a single aggregation point eases the management of subscribers for the ISP, it implies huge bandwidth requirements for the aggregation point and potentially high end-to-end latency for subscribers. An alternative would be a distributed subscriber management, adding more complexity to the management itself. In this paper, a new traffic management architecture is proposed that uses the concept of Software Defined Networking (SDN) to extend the existing Ethernet-based broadband network architecture, enabling a more efficient traffic management for an ISP. By using SDN-enabled home gateways, the ISP can configure traffic flows more dynamically, optimizing throughput in the network, especially for bandwidth-intensive services. Furthermore, a proof-of-concept implementation of the approach is presented to show the general feasibility and study configuration tradeoffs. Analytic considerations and testbed measurements show that the approach scales well with an increasing number of subscriber sessions.

## I. INTRODUCTION

Demand for access to broadband services has grown to an enormous level. In 2012, over 72% of the households in the EU already had broadband connectivity at home [27]. Accordingly, also broadband access architectures are constantly evolving. Since the wider acceptance of ADSL [2], many new access technologies have been introduced, such as xDSL (e.g., ADSL2+, VDSL, and G.SHDSL), optical access, as well as several wireless technologies (e.g., 3G, WiMAX, and LTE). In addition to Internet access services, providers started offering other services to their customers, such as IPTV, Video-on-Demand (VoD), Video Conferencing, Voice over IP (VoIP), and Virtual Private Network (VPN). Furthermore, a multitude of new and innovative cloud-based over-the-top services have been introduced over the years [9]. With the advent of these new services, bandwidth demands have grown tremendously in Internet service provider (ISP) networks.

Traditionally, broadband connections are provided to customers by using Home Gateway (HGW) devices running in

the customer's premises. In the ISP's network, despite having an over-provisioned network core, the access architecture is limited by its design such that all traffic is aggregated to a single point called Broadband Remote Access Server (BRAS), which provides better management options of the subscribers. However, the access to local content and services inside the ISP network is poorly managed within this architecture as all traffic has to pass the BRAS.

In recent years, promising new networking technologies have been introduced in the networking area. In this context, Software Defined Networking (SDN) and OpenFlow [21] as a novel approach for a better network management and virtualization has gained increasing attention. Due to the wide success of IP and fast switching technologies, access networks start migrating towards a network-wide Ethernet support [11], providing the basis for ISP network-wide SDN-based management solutions.

*Goal and Contribution:* In this paper, an SDN-based access network architecture is presented, addressing the above mentioned problems of the single aggregation point for the access to local and ISP-internal services. It is shown how an SDN-enabled HGW device can provide dynamic and flexible configurability to the service provider to optimize traffic flows throughout the network. The conceptual requirements of such an architecture are identified, an experimental design is discussed, and the implemented prototype is presented as well as evaluated as a proof of concept.

The remainder of the paper is structured as follows: Section II provides an overview on related work and background information, Section III introduces the newly proposed architecture, and Section IV discusses its benefits. Section V presents the details on the implemented prototype. Subsequently, Section VI presents the evaluation results and Section VII concludes the paper.

## II. BACKGROUND AND RELATED WORK

The most common technologies used by ISPs for the last mile of the network are Digital Subscriber Line (DSL) and cable, with DSL being the most deployed technology, especially in Europe [25]. The approach presented in this paper therefore primarily targets DSL access networks but should work equally well for cable. The network operator provides connectivity to its subscribers by using the concept of sessions:

<sup>1</sup>Authors contributed equally.

a session has to be established to get access to the network and its services. The session establishment involves procedures for subscriber authentication and application of related network policies. Therefore, the subscriber’s HGW device implements an access protocol to establish a session, which is terminated in the operator’s premises by a device called Broadband Remote Access Server (BRAS)<sup>2</sup>. A typical protocol for this purpose is the Point-to-Point Protocol (PPP), together with its adaptation to work on an Ethernet layer (PPPoE) [12]: all the network packets sent or received by the HGW are encapsulated using PPPoE/PPP headers, creating a point-to-point link with the BRAS. PPP is just one of the options for establishing the subscriber session: different protocols, such as IPoE+DHCP, can be used as well. The HGW is connected to the BRAS through a metropolitan area network via a provider’s access node, called DSLAM. This device establishes the DSL connection with the HGW and provides access to the, e.g., Ethernet-based, aggregation network (see Fig. 1). Usually, the DSLAM uses VLANs to isolate the traffic belonging to different subscribers on the aggregation network. In this architecture, the BRAS implements the subscriber access and policy management, as well as its enforcement and control. Thereby, it centralizes all the functions in a single point, simplifying on the one side network management, but on the other side making the structure of the network rigid. While services could benefit from more flexible structures, they are not easy to support in current operators’ networks. Services like video streaming (e.g. IPTV) could greatly benefit from multicast content delivery within the operator’s network. Unfortunately, current BRASes are the only policy application point for subscriber traffic and maintain point-to-point connections, either because of the use of PPP or because of the VLANs, down to the subscribers’ HGWs. As a result, the aggregation network is logically supporting a set of point-to-point links, limiting the flexibility with which traffic can be routed. Network layer multicast, for instance, cannot use optimized multicast trees when flowing between BRAS and HGWs, since the traffic has to flow through the point-to-point links terminated at the BRAS, taking over the function of the last multicast replication point towards the subscribers. Current approaches to solve such an issue require either to move BRASes towards the edge of the network closer to the subscribers or to build dedicated multicast overlays [10], [11], [13]. Nevertheless, the former goes in a direction contrasting with current research trends, which virtualize BRASes to be run in datacenters [5], while the latter poses challenging management overheads.

### Related Work

The SDN concept is evolving rapidly and already sees first commercial deployments, e.g., in Google’s world-wide backbone connecting its data centers [20] or TeraStream [8], [30], a pilot project by Deutsche Telekom. In this paper, an integrated SDN-based architecture is proposed that includes

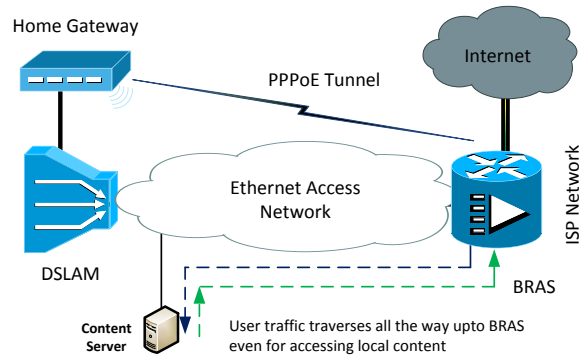


Fig. 1. PPP-based broadband access network architecture

the residential and access network in the network management process to improve access to local and ISP-internal services.

Yiakoumis et al. [31] also propose an SDN-based traffic management approach but with a focus on profile- and subscription-based traffic prioritization within the service provider network. In an earlier work [32], the same authors propose a slicing mechanism at the HGW such that residential gateway hardware is abstracted and virtually sliced among multiple service providers. The authors also propose bandwidth and track isolation between slices using FlowVisor [29]. Berl et al. [3] present an energy efficient virtual home environment which enables resource sharing between users in home networks. The authors also investigate legacy virtualization approaches (i.e. VLANs, VPNs, P2P overlays) as proposed in the VHE Architecture [15], [17], [18]. Abgrall et. al. [1] advance the idea of virtualized HGWs and propose to move functionality away from the home devices and add them to the provider’s access nodes (DSLAMs). This way, the costs for HGW devices could be lowered and the control of the service provider further improved by implementing home network management tasks at the access nodes. A complementary approach is presented by Nick Feamster [14]. The author proposes an SDN-based outsourcing model, where HGW management and operation of home networks is done by a trusted third party entity that has both operations expertise and a broader view on network activities. In contrast to this paper, the above works do not address access to ISP-internal services and its improvement. In most aspects the approaches are orthogonal and complement each other.

Other works in the area investigate the monitoring and logging of network traffic for security and troubleshooting purposes. Calvert et al. [7] propose an OpenFlow-based data recorder tool for HGWs. The authors extend the idea of CWMP (CPE WAN Management Protocol) [6] based on OpenFlow for the use of home network traffic logging. In addition to basic trouble shooting functionality for the service provider, the data recorder tool can provide extended usage statistics. Logged data is stored near the access node and can be shared by the service provider among users for better troubleshooting. In another study Mortier et al. [23] present an OpenFlow-based design for a home router that focuses on

<sup>2</sup>Also called Broadband Network Gateway (BNG)

monitoring and controlling network traffic flows. The proposed idea addresses the desire of home users to better understand and control their network behavior.

### III. SDN-BASED ARCHITECTURE

In this section, we introduce our proposal for a broadband access network architecture. Our proposal evolves the broadband architecture as defined in [12] and briefly described in Section II, in order to allow an easy migration path from current deployments to our solution.

In the proposed architecture, shown in Fig. 2, we substitute the access network between the HGW and the BRAS with a Software Defined Network (SDN). We envision a network composed of SDN-enabled devices, such as OpenFlow switches in the aggregation network, but also SDN-enabled HGWs and BRASes. An SDN controller, using a logically centralized view of the network, manipulates traffic flows according to the operator needs and policies. Exploiting the centralized control view of the SDN controller, the architecture provides a mechanism to free the BRAS from many of the typical functions, such as policy enforcement and traffic monitoring. They can be distributed in the network while still maintaining a simple centralized management point in the controller. Such a distribution allows for a flexible management of subscribers' traffic in the aggregation network, solving one of the main limitations of current architectures, highlighted earlier in this paper.

The two key elements of our architecture are the subscriber session management and the traffic management, which are detailed in the following.

#### A. Subscriber Session Management

The introduction of the subscriber's session concept is a requirement imposed by the need of providing the subscriber with the service he subscribed for and to monitor his activity for several reasons: accounting, to check the compliance with Service-Level Agreements, regulations, network operator policies, and to ensure the correct network operations.

In traditional access architectures, the BRAS is in charge of providing all of the previous functions. We decided to maintain the subscribers' session establishment process at the BRAS also in our architecture. The two main motivations justify this decision: first, we did not intend to define new protocols for establishing subscriber sessions but exploit one of the currently available ones. The BRAS already implements the protocols dedicated to this task, with a complete supporting architecture, e.g., connections to Authentication, Authorization and Accounting (AAA) servers, that have proven to be functional for the purpose; second, using the BRAS for session establishment makes the evolution from current deployments easier, allowing for exploitation of the already deployed hardware.

Among the available protocol options for the session establishment we selected PPPoE<sup>3</sup>, even if we are confident that our approach can be used also when different session management

protocols are in place. Nevertheless, we slightly preferred PPP over the other options (i) for its ability to timely check the session liveness using keep-alive messages and (ii) for its wide deployment in current networks.

In order to get network connectivity, in the first place, the HGW establishes the subscriber session with the BRAS using PPPoE. With the session establishment, the subscriber is assigned with a unique IP address that is used as source network address for all the network flows started by that subscriber, including the flows used for, e.g., IPTV or VoIP services. Note that our architecture does not enforce the usage of publicly routable addresses. In fact, also private IP addresses could be assigned to subscriber inside the operator's network.

Since several functions are moved from the BRAS to the network itself, the network needs to be made aware of the subscriber session and its policies. The introduction of an SDN access network eases this task since it exposes a centralized management point in the SDN controller. Hence, once the BRAS establishes a subscriber session, it informs the SDN controller about the session and its policies. The SDN controller, in turn, can setup the network according to the subscriber related policies.

#### B. Traffic Management

Replacing the Ethernet aggregation network with an SDN provides high flexibility in the way flows routing strategies are taken. In a traditional architecture, the DSLAM has to rely on VLANs to ensure the proper traffic steering towards the BRAS. In particular, the DSLAM applies one or more VLAN tags in order to identify the subscriber line and the related service (e.g. VoIP). The BRAS removes the VLAN tags after using them to identify traffic flows and apply proper policies, such as traffic shaping. VLANs are a strict requirement in traditional networks, since otherwise the Ethernet aggregation network would provide a single broadcast domain to all the subscribers connected to it. In an SDN, instead, the actual flow forwarding is defined by the SDN controller that is adding/deleting/modifying flow forwarding rules at the network devices, enabling us to define, from scratch and out of any legacy constraint, the way subscriber flows traverse the network up to the BRAS.

Keeping in mind this shift in the network operations, we have to address three main classes of network flows of a subscriber:

- 1) Session establishment: network flows related to the establishment of the subscriber session;
- 2) Internal traffic: any network flow between a subscriber and an end-point located on the same access network, i.e. before the BRAS;
- 3) External traffic: any network flow between a subscriber and an end-point located on the Internet, i.e. behind the BRAS.

For the session establishment, each subscriber needs to be able to exchange PPPoE packets with the BRAS. In the proposed architecture, each subscriber has a dedicated path to the BRAS for the PPPoE traffic, including both the traffic needed

<sup>3</sup>In this paper we refer to the combination of PPP/PPPoE using simply the PPPoE acronym for brevity.




Fig. 2. SDN-based traffic management architecture

for the session establishment and the traffic belonging to the PPPoE tunnel created between the HGW and the BRAS once the session is established. Since in traditional architectures any packet sourced/destined from/to a subscriber is sent using the PPPoE tunnel, this approach facilitates compatibility with legacy architectures. That is, if the SDN controller configures only forwarding rules related to the session establishment flows, the proposed architecture falls back to the current legacy architectural approaches. The session establishment flow is crucial in the proposed architecture, since any other subscriber's network flow is closely bound to it. In fact, any other subscriber's flow in the network is allowed by the SDN controller only if the subscriber session has been established.

For the handling of internal traffic, the approach exploits the full flexibility of SDN. In fact, having PPPoE in place, in a traditional network, the traffic handled on the WAN interface of an HGW should be encapsulated in a PPPoE tunnel. This clearly requires a decapsulation/encapsulation function on both end-points of the tunnel. If the HGW implements this function, it is also to be implemented at the BRAS. With the proposed architecture, not all the traffic is directed to the BRAS anymore, since we want to enable, e.g., local traffic offloading for service break-outs. To account for this issue, we consider also HGWs to be part of the devices managed by the SDN controller. In particular, we assume that HGWs can provide a configurable decapsulation/encapsulation function that allows for a dynamic, fine grained decision on the network flows that are to be encapsulated in the PPPoE tunnel. Exploiting this function, the SDN controller can instruct a HGW to not encapsulate a given network flow, hence, a plain IPoE network flow would be sent out from the WAN interface of the HGW. The SDN controller can then define an arbitrary route for this network flow, for instance, to steer the traffic towards a local content cache server.

For external traffic, the SDN controller can choose the PPPoE tunnel, the solution adopted for internal traffic, or a mix of both, by configuring the HGW encapsulation function.

A sensible point for the functional correctness of the proposed traffic management approach is the ability to still enforce the correct operator's policies and monitoring actions to the subscriber network flows, in case such flows are not passing through the BRAS. Once more, we exploit SDN devices deployed in the network to provide both policy enforcement and monitoring functionality. Policy enforcement, such as Access Control Lists (ACL) and traffic shaping, can be provided using switch and/or SDN-enabled DSLAM capabilities. Monitoring in terms of number of bytes or connection times can be inferred from the flow counters installed at the SDN devices, provided that those counters are differentiated per subscriber (and also per service if needed). Assigning unique IP addresses to subscribers allows simplifying the network flow identification for the SDN controller, in particular when the policy enforcement and monitoring actions need to be distributed across several network devices.

### C. A day in the life of a subscriber session

We now describe an example of operations in the case of a subscriber session establishment, assuming that the network provides, on top of an Internet access service, additional services such as VPN and access to local content servers.

The subscriber's HGW starts the PPPoE session establishment with the BRAS by using the pre-configured path between the DSLAM to which the HGW is connected and the BRAS. Once the session is established, the BRAS notifies the SDN controller and provides it with the subscriber's policies details. In turn, the SDN controller installs the rules to implement the subscriber's ACL and traffic shaping policy on the SDN-enabled DSLAM, while the access network is

configured according to the flows routing strategy. In our example, reported in Fig. 2, apart from the session establishment flows, depicted using the “PPPoE tunnel”, two more paths are installed in the network for the subscriber: (i) an IPoE flow based on, e.g., destination and source IP addresses, is steered to another subscriber’s premises to implement a VPN between the two subscribers, while (ii) a second IPoE flow is directed to a content server, that contains, e.g., the movies the subscriber bought using a video-on-demand service offered by the network operator.

In parallel with the configurations performed by the SDN controller on the network, the SDN-enabled HGW, right after the PPPoE session establishment, establishes a control session with the SDN controller in order to download the PPPoE encapsulation policy, so that the paths configured in the network can be correctly accessed by the traffic generated at the HGW.

#### IV. BENEFITS OF SDN-BASED ARCHITECTURE

The introduction of an SDN-based architecture provides several benefits over current approaches: first, it allows for an efficient utilization of the resources within the network and helps in achieving low latency and high throughput for time critical services such as IPTV, without having to compromise on security or ease of management in comparison to existing PPP-based network architectures; second, access to local content will not require additional AAA mechanism as proposed in some other architectures [28]; finally, we can have flexible and granular control over traffic flows throughout the network, with the ability to implement in-network Quality of Service (QoS) and traffic policies right at the edge of the network, without requiring a single aggregation point for policy enforcement. The previous list names just a few of the benefits the SDN-based architecture would introduce. In the remainder of this section, these benefits are further motivated by briefly describing some interesting use cases.

##### A. Peer-to-Peer Services within the ISP Network

Consider a scenario where one subscriber is trying to access content from another subscriber within the same service provider network. Let us assume that both subscribers are geographically co-located (terminated at the DSLAM or co-located neighboring DSLAMs). In a traditional broadband network, user traffic would be encapsulated as a PPPoE tunnel and traverse all the way up to the BRAS.

Now consider the proposed architecture. In this scenario, user traffic is routed directly from one DSLAM to the neighboring DSLAM as an IPoE connection. By doing so, we can eliminate excessive traffic congestion inside the core network, while at the same time provisioning subscriber access using PPP-based authentication at the BRAS. Moreover, the SDN controller can combine the path selection strategy with the information provided to P2P applications using ALTO (Application-Layer Traffic Optimization) like services [26].

##### B. Multicast Traffic using SDN

One of the problems with a PPPoE-based architecture is having access to multicast traffic stream for individual subscribers. Since every subscriber’s PPPoE tunnel endpoint is at the BRAS, after the BRAS the traffic has to be distributed as a unicast stream, despite the fact that multiple subscribers accessing this stream could be co-located at the same DSLAM. For the specific case of IPTV, there have been new approaches defined such as in [28], where a separate security and management mechanism is provisioned for such services. In existing architectures, solutions are usually based on the provisioning of dedicated multicast VLANs managed by distributing IGMP snooping functions in the devices along the network (both DSLAMs and/or BRASes). Those VLANs usually require a separated management from the BRAS for the enforcement of the policies, introducing further complexity in the architecture.

In our proposal, instead, multicast traffic can be detected, e.g., by instructing the switches on the path to forward IGMP packets towards the SDN controller. The controller, in turn, can reconfigure the network implementing a proper multicast distribution. Therefore, in case of such a multicast traffic, the SDN controller can configure the network such that a single unicast stream can be forwarded till the last distribution point (e.g. the DSLAM on which multiple users are connected accessing the same streaming content) from where traffic can be distributed to individual subscribers.

#### V. PROOF OF CONCEPT

This section presents the prototype implementation as a proof of concept for the proposed architecture.

##### A. Home Gateway

The proposed SDN extensions have been added to the HGW in the form of separate modules based on the OSGi platform<sup>4</sup>, which packages each piece of code into bundles, i.e. functional entities that can be deployed dynamically. The implemented HGW OSGi bundle, depicted in Fig. 3, attempts to set up a PPPoE connection with the BRAS using the PPP configurator. Once the connection is up, the PPP configurator informs the routing module about the connection state. The routing module works in tandem with the HGW-to-SDN communicator, which requests the details about the flows to be routed directly over the IPoE connection from the SDN controller. The communicator then passes these details to the routing module to configure the IP routing.

The communication between the HGW and the SDN controller happens through an in-band channel that uses the PPPoE tunnel established with the BRAS. Thus, the HGW-to-SDN communicator sends out a configuration request, using a stateless HTTP-based protocol<sup>5</sup> through the same PPP interface that is used for the subscriber connection to the BRAS. The HTTP communication can be encrypted using SSL in order to achieve a secure communication between the SDN controller and the HGW.

<sup>4</sup><http://www.osgi.org/>

<sup>5</sup>The SDN controller implements a REST interface for this purpose.




Fig. 3. Components of the SDN controller and the HGW OSGi bundle

### B. BRAS

In order to implement the proposed architecture, our design for the BRAS consists of a device that supports two functions for the realization of our prototype: (i) the PPP session termination and (ii) the BRAS to SDN controller communication for per user session state updates. Correspondingly, the BRAS device consists of two components, namely the PPP server and the BRAS-to-SDN controller communicator. The PPP server updates the BRAS-to-SDN controller communicator with the session state of each connected user and forwards it to the SDN controller.

In our prototype, we implemented the PPP session termination using both a GNU/Linux server running a customized version of the RP-PPPoE<sup>6</sup> software and a high performance virtual BRAS implementation presented in [5].

### C. Controller

The implementation of the SDN controller extends an existing OpenFlow-based controller that was modified to enable the interaction with the external modules managing the communication with the HGWs and the BRASes (see Fig. 3).

The **SDN application** runs within the OpenFlow controller framework, modifying its standard routing behavior. It is used to configure and modify a network's flow paths by interacting with OpenFlow-enabled devices.

The **HGW listener** is a standalone HTTP server module. Upon receiving a request from a particular subscriber, the HGW listener extracts its unique IP address and retrieves the encapsulation policies to apply for this particular subscriber, which is then passed to the HGW as an HTTP response.

The **BRAS listener** is a TCP server that listens to incoming messages from the BRAS. Whenever a PPPoE connection is set up or terminated, the BRAS sends a subscriber session state update to the BRAS listener. The BRAS listener extracts the subscriber IP address, connection state information and subscriber policies. This information is then passed to the SDN Application which configures the network devices accordingly.

An alternative approach to the BRAS listener would have been to intercept PPPoE discovery handshake messages,

namely PADI and PADT, flowing through the network. With this information the SDN controller could learn the connection state of individual subscribers in advance. However, considering a large network, this approach would cause scalability concerns for the increased amount of signaling to be handled at the SDN controller. Also, the SDN controller would still require to retrieve the subscriber policy, e.g., from the BRAS. Therefore, we opted to use explicit session update signaling using the BRAS listener.

## VI. EVALUATION AND DISCUSSION

In Section V we presented an implemented prototype of the proposed architecture. We evaluated our prototype using two different testbed setups.

A first physical testbed, shown in Fig. 4, has been used in order to evaluate the achieved session establishment rate when configuring the routing strategy, as implemented at the SDN controller, on a network of OpenFlow switches. The testbed uses a NEC PF5240 switch [24], configured to implement four virtual switch instances; a PPP session generator has been developed as well, in order to support a session generation rate up to 1000 sessions per second. For this test we adopted the BRAS presented in [5].

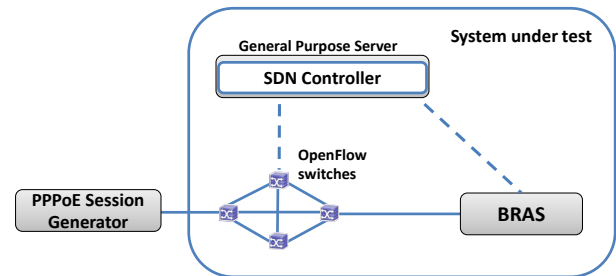


Fig. 4. Testbed used for the PPP session establishment rate evaluation

A second virtual testbed, reproducing the setup shown in Fig. 3, is composed of five virtual machines, which emulate the access network, and of a physical machine, which hosts the SDN controller. The virtual machines implement, respectively, the HGW, a DSLAM, an OpenFlow network composed by four switches (we used Mininet [16] to this purpose), a content cache server and a BRAS (implemented using RP-PPPoE).

<sup>6</sup><http://www.roaringpenguin.com/products/pppoe>



We added OpenFlow functions to the DSLAM by cascading it with an OpenFlow software switch.

In the remainder of this section, we present the evaluation of the prototype that we performed using the described testbeds.

### A. Session Establishment Rate

Using the physical testbed, we successfully generated and established up to 1000 sessions per second. In particular, during the test we increased the generation rate at the PPP session generator, starting from 100 up to 1000 sessions per second. Sessions are generated uniformly in one second. Fig. 5 shows that our prototype was able to sustain the offered load in all the cases.

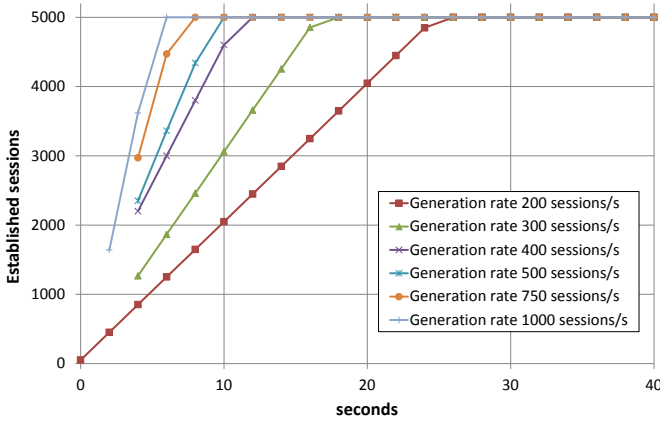


Fig. 5. Achieved session establishment rates

### B. SDN Controller Scalability

An important element that affects the scalability in SDN-based networks is the software controller, whose scalability depends on the ability to simultaneously handle a large number of switching devices. To address this challenge by design, we used a modular controller architecture similar to the one presented in [4]. Besides, the controller can exploit well-understood mechanisms used to build scalable software systems, such as clustering, replication, or distribution of the components. While the overall scalability of the controller design certainly is an important issue, in the following, we focus on the most critical module specific to the presented approach: the HGW listener. This module could impact the scalability of the controller since it has to handle the connections of a potentially high number of HGWs.

Since, depending on the characteristics of the local services to be supported, the communication between controller and HGW might not be time-critical, we decided to use a polling-based communication for this purpose. Alternatively to the polling-based mechanism, a permanent connection between the HGWs and the controller could be used to inform the HGWs about changes in the routing policies. While this would allow instant and system-wide updates, it clearly does not scale well with a growing number of subscribers. To better understand the trade-off between a polling-based approach and an always-on connection, Fig. 6 exemplarily shows the number

of active connections for the two cases. For the polling-based approach, two parameters play an important role: the polling interval and the duration of an individual connection. For the plots, a total number of 1000 users is assumed that establish sessions at a random point in time between 0 and 60 seconds.

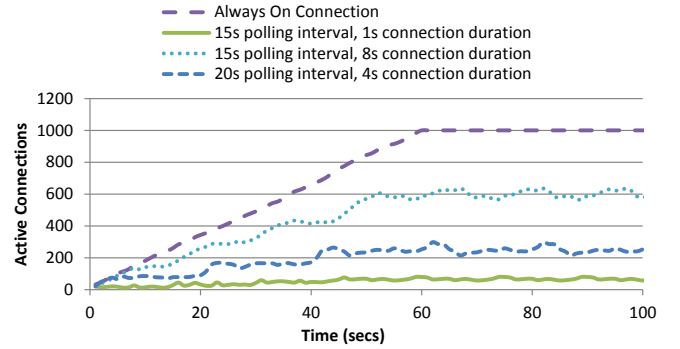


Fig. 6. Comparison of polling-based and always-on HGW reconfiguration with different polling intervals and connection durations

The analytical and exemplary results depicted in Fig. 6 show the expected advantage of a polling-based mechanism in that it reduces the number of synchronously active connections to the controller, compared to the always-on alternative. It further shows how the number of active connections decreases for an increase in the polling interval. Similarly, depending on the connection payload and SDN controller RTT, the connection duration of individual polls affects the performance such that with larger connection durations, there are more active connections at the controller at a given point in time. In realistic scenarios, we expect connection durations and RTTs below the second.

### C. SDN Devices Scalability

In the proposed design, one of the key advantages of having an SDN-based setup is the ability to re-route subscriber traffic dynamically. Nevertheless, there are studies (e.g., [19], [22]) showing that, depending on the device, there could be important limitations in the ability of the switch to handle high rates in forwarding table updates. In our case, the issue could be a limiting factor in particular at the edge of the network, i.e. at the DSLAM (or at the SDN switch co-located with the DSLAM as in our virtual testbed), where per subscriber forwarding rules, such as the ones required to implement ACLs, are needed. In the other parts of the network, we expect that the issue can be alleviated if flows aggregation strategies are used.

In order to check the ability of our architecture to meet current SDN devices' limitations, we need to evaluate the required rule update rate at a DSLAM. Assuming a DSLAM to be connected to a maximum of 1000 subscribers, even in the unrealistic case in which all the subscribers establish a session at the same time, i.e. requiring the installation of forwarding rules at the DSLAM, the amount of rules to be installed would be still manageable. For instance, assuming that 10 rules per subscriber are required, and that the DSLAM can sustain a rate

of 50 rules per second, it would take 200 seconds to install all the rules. Moreover, our architecture is able to already provide the traditional connectivity service by solely installing the flows related to the session establishment (see Section III), which requires just a single rule per subscriber, reducing the total required amount of rules to be installed to 1000 and the time needed to handle them to 20 seconds.

## VII. CONCLUSION AND FUTURE WORK

In this paper we propose an evolution to the current broadband network architecture based on SDN in order to flexibly support advanced services, such as the provisioning of content-based services like IPTV, while still relying on a BRAS-based subscriber session management. By re-routing traffic of authenticated subscribers for local and ISP-internal services already inside the access network, an optimized service access can be realized while removing costly load from the BRAS. To achieve this purpose, it was shown how existing BRAS-based session management can be integrated with an OpenFlow-based access network and a custom HGW to SDN controller communication. An evaluation of the system and the analysis of key scalability issues show that our approach is applicable even to networks with large user bases. In our future work, we aim to evaluate additional aspects of our approach including performance and cost trade-offs. Furthermore, several security aspects, which are not part of this work, such as the implications of communication channels between HGW and SDN controller, should be taken into account.

## ACKNOWLEDGMENT

This work has been funded in parts by the European Union (FP7/#317846, SmartenIT and FP7/#318398, eCousin) and the DFG as part of the CRC 1053 MAKI.

## REFERENCES

- [1] D. Abgrall, "Virtual Home Gateway, How can Home Gateway virtualization be achieved?" Eurescom, Tech. Rep. P2055, 2011, <http://archive.eurescom.eu/~pub/deliverables/documents/P2000-series/P2055/D1/P2055-D1.pdf>.
- [2] American National Standards Institute, "Network and Customer Installation Interfaces Asymmetric Digital Subscriber Line (ADSL) Metallic Interface," 1998, ANSI T1.413.
- [3] A. Berl, R. Weidlich, M. Schrank, H. Hlavacs, and H. de Meer, "Network Virtualization in Future Home Environments," in *IFIP/IEEE DSOM*, 2009.
- [4] R. Bifulco, R. Canonico, M. Brunner, P. Hasselmeyer, and F. Mir, "A Practical Experience in Designing an OpenFlow Controller," in *IEEE EWSDN*, 2012.
- [5] R. Bifulco, T. Dietz, F. Huici, M. Ahmed, J. Martins, S. Niccolini, and H.-J. Kolbe, "Rethinking Access Networks with High Performance Virtual Software BRASes," in *IEEE EWSDN*, 2013.
- [6] BroadbandHome Working Group, "CPE WAN Management Protocol," Broadband Forum, Tech. Rep. No. 69, Issue 1, Amendment 4, 2011, [http://www.broadband-forum.org/technical/download/TR-069\\_Amendment-4.pdf](http://www.broadband-forum.org/technical/download/TR-069_Amendment-4.pdf).
- [7] K. L. Calvert, W. K. Edwards, N. Feamster, R. E. Grinter, and Y. D. X. Zhou, "Instrumenting Home Networks," *SIGCOMM Computer Communication Review*, vol. 41, no. 1, pp. 84–89, 2011.
- [8] Deutsche Telekom, "Deutsche Telekom tests TeraStream, The Network of the Future, in Croatia," <http://www.telekom.com/media/company/168008>, [Accessed 2013-08-29].
- [9] Earnest & Young, "Cloud Computing Issues and Impacts," Tech. Rep. EYG No DC0078, 2011, [http://www.ey.com/Publication/vwLUAssets/Cloud-computing\\_issues\\_and\\_impacts/\\$FILE/Cloud\\_computing\\_issues\\_and\\_impacts.pdf](http://www.ey.com/Publication/vwLUAssets/Cloud-computing_issues_and_impacts/$FILE/Cloud_computing_issues_and_impacts.pdf).
- [10] End to End Architecture Working Group, "DSL Evolution - Architecture Requirements for the Support of QoS-Enabled IP Services," Broadband Forum, Tech. Rep. 59, 2003, <http://www.broadband-forum.org/technical/download/TR-059.pdf>.
- [11] —, "Migration to Ethernet-Based DSL Aggregation," Broadband Forum, Tech. Rep. 101, 2006, <http://www.broadband-forum.org/technical/download/TR-101.pdf>.
- [12] —, "Migration to Ethernet-based DSL Aggregation - Issue 2," Broadband Forum, Tech. Rep. No. 101, Issue 2, 2011, [http://www.broadband-forum.org/technical/download/TR-101\\_Issue-2.pdf](http://www.broadband-forum.org/technical/download/TR-101_Issue-2.pdf).
- [13] —, "Multi-service Broadband Network Functional Modules and Architecture," Broadband Forum, Tech. Rep. 145, 2012, <http://www.broadband-forum.org/technical/download/TR-145.pdf>.
- [14] N. Feamster, "Outsourcing Home Network Security," in *ACM SIGCOMM HomeNets Workshop*, 2010.
- [15] A. E. Garcia, A. Berl, K. A. Hummel *et al.*, "An Economical Cost Model for fair resource sharing in Virtual Home Environments," in *IEEE Next Generation Internet Networks*, 2008.
- [16] N. Handigol, B. Heller, V. Jeyakumar, B. Lantz, and N. McKeown, "Reproducible Network Experiments Using Container-based Emulation," in *ACM CoNEXT*, 2012.
- [17] H. Hlavacs, R. Weidlich, K. A. Hummel *et al.*, "Distributed Energy Efficiency in Future Home Environments," *Annals of Telecommunications*, vol. 63, no. 9-10, pp. 473–485, 2008.
- [18] H. Hlavacs, R. Weidlich, and T. Treutner, "Energy Saving in Future Home Environments," in *Home Networking Conference at IFIP Wireless Days*, 2008.
- [19] D. Y. Huang, K. Yocum, and A. C. Snoeren, "High-fidelity Switch Models for Software-Defined Network Emulation," in *ACM SIGCOMM HotSDN Workshop*, 2013.
- [20] S. Jain, A. Kumar, S. Mandal, J. Ong, L. Poutievski, A. Singh *et al.*, "B4: Experience with a Globally-deployed Software Defined WAN," in *ACM SIGCOMM*, 2013.
- [21] N. McKeown, T. Anderson, H. Balakrishnan, G. Parulkar *et al.*, "OpenFlow: Enabling Innovation in Campus Networks," *SIGCOMM Computer Communication Review*, vol. 38, no. 2, pp. 69–74, 2008.
- [22] J. C. Mogul, J. Tourrilhes, P. Yalagandula, P. Sharma, A. R. Curtis, and S. Banerjee, "DevoFlow: Cost-effective Flow Management for High Performance Enterprise Networks," in *ACM SIGCOMM HotNets Workshop*, 2010.
- [23] R. Mortier, T. Rodden, T. Lodge, D. McAuley, C. Rotsos, A. Moore, A. Koliouisis, and J. Sventek, "Control and Understanding: Owning your Home Network," in *IEEE COMSNETS*, 2012.
- [24] NEC, "ProgrammableFlow PF5240 Switch," <http://www.necam.com/docs/?id=5ce9b8d9-e3f3-41de-a5c2-6bd7c9b37246>.
- [25] OECD, "Communications Outlook 2011 – Table 4.16. Total Broadband Subscriptions by Access Technology," <http://dx.doi.org/10.1787/888932398138>.
- [26] J. Seedorf and E. Burger, "Application-Layer Traffic Optimization (ALTO) Problem Statement," RFC 5693 (Informational), Internet Engineering Task Force, Oct. 2009, <http://www.ietf.org/rfc/rfc5693.txt>.
- [27] H. Seybert, "EuroStat Statistics in Focus 50/2012," Tech. Rep. Catalogue number KS-SF-12-050-EN-N, 2012, [http://epp.eurostat.ec.europa.eu/cache/ITY\\_OFFPUB/KS-SF-12-050/EN/KS-SF-12-050-EN.PDF](http://epp.eurostat.ec.europa.eu/cache/ITY_OFFPUB/KS-SF-12-050/EN/KS-SF-12-050-EN.PDF).
- [28] C. Sherbo, "Broadband Network Architecture Options for the Delivery of IPTV," 2007, Presentation at NAB, Las Vegas.
- [29] R. Sherwood, G. Gibb, K.-K. Yap, G. Appenzeller, M. Casado, N. McKeown, and G. Parulkar, "FlowVisor: A Network Virtualization Layer," Openflow Switch Consortium, Tech. Rep., 2009.
- [30] Tail-f Systems, "Deutsche Telekom Selects Tail-f as Provider of Software-Defined Networking in TeraStream Project," <http://finance.yahoo.com/news/deutsche-telekom-selects-tail-f-120000936.html>, [Accessed 2013-08-29].
- [31] Y. Yiakoumis, S. Katti, T.-Y. Huang, N. McKeown, K.-K. Yap, and R. Johari, "Putting Home Users in Charge of their Network," in *ACM Conference on Ubiquitous Computing*, 2012.
- [32] Y. Yiakoumis, K.-K. Yap, S. Katti, G. Parulkar, and N. McKeown, "Slicing Home Networks," in *ACM SIGCOMM HomeNets Workshop*, 2011.