# Software-Defined Crowd-Shared Wireless Mesh Networks

Ahmed Abujoda[†] Arjuna Sathiaseelan[§] Amr Rizk[†] Panagiotis Papadimitriou[†]

[†]Institute of Communications Technology, Leibniz Universität Hannover, Germany
{*ahmed.abujoda, amr.rizk, panagiotis.papadimitriou*}*@ikt.uni-hannover.de*
[§]Computer Laboratory, University of Cambridge, UK
*arjuna.sathiaseelan@cl.cam.ac.uk*

***Abstract*—Universal access to Internet is crucial, and as such, there have been several initiatives to enable wider access to the Internet. Public Access WiFi Service (PAWS) is one such initiative that takes advantage of the the available unused capacity in home broadband connections and allows Less-than-Best Effort (LBE) access to these resources, as advocated by Lowest Cost Denominator Networking (LCDNet). PAWS has been recently deployed in a deprived community in Nottingham, and, as any crowd-shared network, it faces limited coverage, since there is a single point of Internet access per guest user, whose availability depends on user sharing policies.**

**To mitigate this problem and extend the coverage, we consider a crowd-shared wireless mesh network (WMN) in which the home routers are interconnected as a mesh. Such a network provides multiple points of Internet access and can enable resource pooling across all available paths to the Internet backhaul. In this paper, we investigate the potential benefits of a crowd-shared WMN for public Internet access by performing a comparative study between such a network and PAWS. To this end, we present a software-defined WMN control plane for the coordination of traffic redirections through the WMN and an algorithm for Internet access point selection. Our simulation results show that a crowd-shared WMN can provide much higher utilization of the shared bandwidth and can accommodate a substantially larger volume of guest user traffic.**

## I. INTRODUCTION

The Internet has evolved into a critical infrastructure for education, employment, e-governance, remote health care, digital economy, and social media. However, the Internet today is facing the challenge of a growing digital divide, i.e., an increasing disparity between those with and without Internet access [20]. Access problems often stem from sparsely spread populations living in physically remote locations, since it is simply not cost-effective for Internet Service Providers (ISPs) to deploy the required infrastructure for broadband Internet access in these areas. Coupled with physical limitations of terrestrial infrastructures (mainly due to distance) to provide last mile access, remote communities also incur higher costs for connection between the exchange and backbone network when using wired technologies, because the distances are longer. Ubiquitous mobile broadband coverage is currently not feasible, since direct investment in local infrastructure is uneconomic [15]. Addressing digital exclusion due to socio-economic barriers is also important. The United Nations revealed the global disparity in fixed broadband access, showing that access to fixed broadband in some countries costs almost 40 to 100 times their national average income [8].

The reluctance of network operators (who are economically motivated) to provide wired and cellular infrastructures to rural/remote areas have led to several initiatives to build large-scale, self-organized, and decentralized community wireless networks that use WiFi mesh technology (including long distance), due to the reduced cost of using the unlicensed spectrum. These community wireless mesh networks have self-sustainable business models, which provide more localised communication services, as well as Internet backhaul support via peering agreements with traditional network operators who see such networks as a way to extend their reach at a lower cost. There also are community-led wireless initiatives such as crowd-shared wireless networks, in which home broadband owners share a portion of their home broadband with friends, neighbours, or other users either for free or as part of a service offering by the ISP.

Public Access WiFi Service (PAWS) [13] is a community-led crowd-shared WiFi service that uses a set of techniques that make use of the available unused capacity in home broadband networks and allowing Less-than-Best Effort (LBE) access to these resources [18]. PAWS adopts an approach of community-wide participation, where home broadband subscribers are enabled to donate controlled but free use of their high-speed broadband Internet to fellow citizens.

PAWS is currently under deployment with 20 custom-made PAWS routers placed in a deprived community in Nottingham and is also being trialled out in rural Wales. PAWS currently serves as a medium-scale open network measurement observatory in the UK allowing researchers to gather data about network availability, reachability, topology, security, and broadband performance from distributed vantage points in socio-economically deprived urban and rural areas. The PAWS deployment is essentially a crowd-shared access network (similar to FON) for under-privileged users in urban and rural communities. This provides the research community with a wealth of information on the needs of under-privileged users in terms of their access patterns and what do they use Internet access for. However, PAWS has faced ongoing deployment challenges, such as limited coverage, due to home user sharing patterns (i.e., home network users stop sharing their Internet connection for certain periods).

In this paper, we investigate the potential benefits of enabling PAWS or any crowd-shared wireless network as a crowd-shared wireless mesh network (WMN). We particularly consider crowd-shared WMNs in residential areas, where the dense deployment of wireless home routers allows their interconnection as a mesh. To take advantage of the multiple points of Internet access in such a WMN, we are currently deploying a software-defined WMN (SDWMN) control plane in one of the CONFINE community networks [1] for the redirection of guest user traffic to any of the available access points. Since traffic redirection depends on user sharing policies, we have generated a model for home network user sharing patterns based on the router on/off periods captured from the PAWS network. Using simulations, we quantify the benefits of a crowd-shared WMN for Internet access sharing. Specifically, we show that a WMN can achieve very high utilization of the shared bandwidth and can accommodate a significantly larger volume of guest user traffic compared to a crowd-shared network with a single point of access, such as PAWS.

The remainder of the paper is organized as follows. In Section II, we present the design of a SDWMN control plane and its deployment in a CONFINE community network. Section III discusses techniques for guest user traffic redirection. In Section IV, we present our simulation environment and discuss the benefits of using a crowd-shared WMN for public Internet access. Section V briefly discusses related work. Finally, Section VI highlights our conclusions.

## II. Software Defined Crowd-Shared Wireless Mesh Networks

The underlying problem with PAWS or any crowd-shared network is that they serve as single point of Internet access to guest users within the coverage of the wireless router and hence, they have no provision to extend the coverage when no bandwidth is being shared. Based on our experience from the trial PAWS deployment, PAWS routers were not available for certain periods, because sharers needed all the bandwidth of their broadband connection or due to other reasons, such as economic constraints placed on home users in underprivileged areas where they are enforced to conserve energy by turning off the routers at nights. These observed user behaviors entail significant challenges for the successful adoption of PAWS.

### *A. Crowd-Shared WMN Management*

A potential solution to this problem is to extend the PAWS network as a crowd-shared WMN. Such a network would allow home network users to share part of their own broadband connection to the public for free while also connected to each other as a WMN providing extended coverage (Fig. 1). Extending PAWS to a crowd-shared WMN departs from the norm: multiple users from different ISPs form part of the WMN to provide free Internet connectivity, while most wireless community WMNs today are operated by a single organization. This raises important questions regarding the operation, configuration, and management of crowd-shared WMNs.

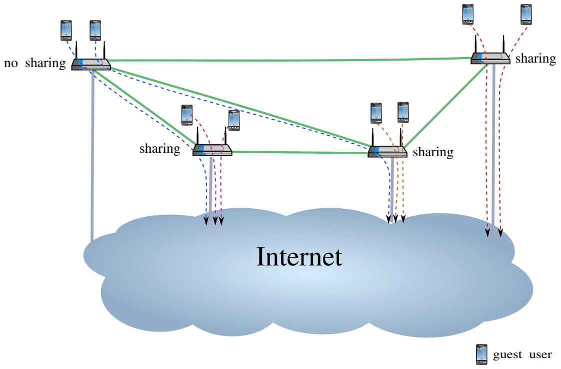

Fig. 1. Crowd-shared WMN for public Internet access.

Software-defined networking (SDN) can facilitate the management and operation of wireless networks at large scale. Leveraging on SDN's centralized control and network-wide visibility, the management and operation of a crowd-shared WMN can be outsourced to a third party. In [14], we describe a holistic approach of coupling both social and economic incentives in designing future networks allowing the extension of the stakeholder value chain to include more than the two traditional parties (consumer and Internet service provider). Compared to existing mesh networks for Internet access sharing (e.g., Freifunk [3]), our approach provides more opportunities for non-governmental organizations and local governments (driven by social goals rather than economic) to become virtual network operators. Enabling a third party to federate such wireless home networks would reduce the operating expenditures for network operators as well as enable new economic models for revenue generation from currently underutilized infrastructures. In particular, we rely on SDN to create the notion of Virtual Public Networks (VPuN), i.e., crowd-shared home networks created, deployed and managed through an evolutionary SDN control abstraction [14]. Although originally intended for crowd-shared wireless networks such as PAWS, VPuN can be also used for crowd-shared WMNs, enabling resource pooling across multiple home broadband connections based on the prevailing network conditions and usage sharing patterns.

### *B. Control Plane Deployment*

Based on the notion of VPuN, we are currently deploying a SDWMN control plane in one of the CONFINE community networks. These community networks consist of research devices (RDs) that host isolated containers (i.e., so-called slivers) which can be controlled and configured by a user, according to the needs of his experiment [5]. The CONFINE network used for the deployment of our control plane is a wireless mesh with around 30 research devices and more than 1000 wireless nodes (i.e., each RD is connected to the WMN).

Fig. 2 illustrates the experimental setup and an overview

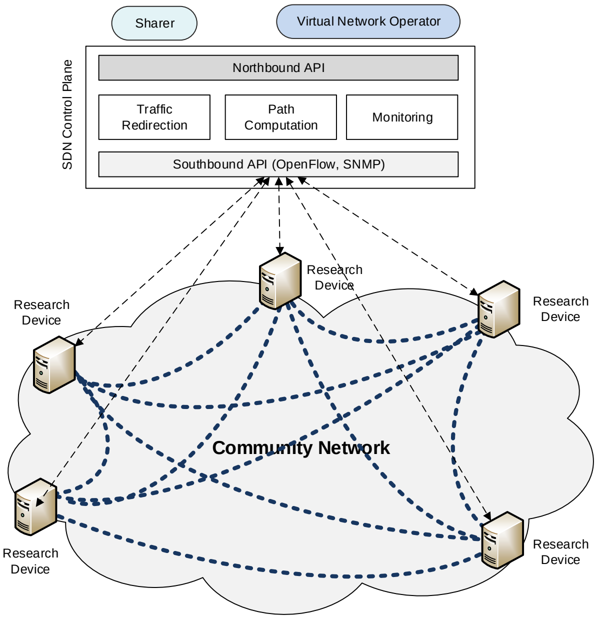

Fig. 2. Software-defined WMN control plane overview.

of the SDWMN control plane. The main goal of the control plane is to improve the utilization of shared bandwidth by redirecting guest user traffic to one of the home routers through the WMN, based on user sharing policies. As such, the control plane exposes an interface to the home network user allowing him to express a sharing policy (e.g., using the sharing policy expression language presented in [14]). We briefly discuss only the main control plane modules, which include path computation, traffic redirection, and monitoring. Path computation is responsible for computing shortest paths between a pair of RDs. The traffic redirection module selects the gateway for Internet access (using the algorithm presented in III) and subsequently sets up a tunnel[1] between the RD where the guest user is connected and the RD assigned as the gateway (this procedure is explained in detail below). The monitoring module collects statistics about the bandwidth utilization of the WMN and the access links.

Traffic redirection is carried out in a sliver using the following components (Fig. 3):

- An encapsulation module, implemented with Click Moduler Router [10], encapsulates all incoming traffic using IP-in-UDP. The destination IP address appended to the tunnel header corresponds to the assigned gateway IP address. A separate encapsulation module is set up for each tunnel.
- A decapsulation module, implemented with Click, strips the tunnel header and extracts the original packet.
- A management module instantiates or terminates an encapsulation module when a tunnel needs to be set up or torn down, respectively. Such actions are triggered by messages received from the SDWMN control plane.
- A virtual switching module, implemented with Open-vSwitch [12], steers data traffic between the physical ports and the modules for encapsulation/decapsulation, and also forwards messages from the SDWMN control plane to the management module.

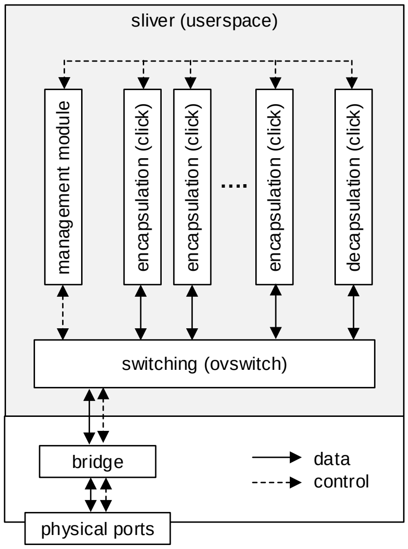

Fig. 3. Traffic redirection in a CONFINE research device.

For each new flow, the control plane assigns a gateway, instructs the management module to install a new encapsulation module (if a tunnel to the assigned gateway has not been previously set up), and inserts the required entries to the flow table of the virtual switch so that this flow is sent to the corresponding encapsulation module. Guest user flows are identified using the source/destination IP addresses and ports.

## III. Traffic Redirection

To take advantage of the extended coverage provided by a WMN, we develop and evaluate techniques for the redirection of guest user traffic during the periods that the home user does not permit the sharing of his broadband connection. In this case, Internet is accessed through the router of another home network where sharing is allowed.

In the following, we present an algorithm (Algorithm 1) for the assignment of the gateway and the path over which the traffic will be redirected through the WMN. This algorithm will be executed by a SDN controller which has knowledge of the WMN topology and utilization, the Internet access links utilization, and the user sharing policies.

We represent the WMN as a weighted undirected graph $G = (N, L)$, where $N$ is the set of nodes and $L$ is the set of links between nodes of the set $N$. Each node $n_i$ is associated with an Internet access link whose available shared bandwidth is denoted by $C(n_i)$. Each link $l_{ij} \in L$ between two nodes $n_i$ and $n_j$ is associated with the available bandwidth $C(l_{ij})$. Let $P_{ij}$ denote the set of paths in the network $G$, between the pair of nodes $n_i$ and $n_j$. The available bandwidth $C(p)$ associated to

[1] Since the wireless nodes in the community network are not under our control, we redirect traffic through tunnels set up between pairs of slivers. A tunnel is dynamically set up when traffic has to be redirected over that path, and it is torn down when it is no longer needed.

a path $p \in P_{ij}$ is given by the minimum residual bandwidth of the links along the path:

$$C(p) = \min_{l_{ij} \in p} C(l_{ij}) \quad (1)$$

We further represent a flow demand for a guest user with the tuple $d = (n_u, r)$, where $n_u \in N$ is the node where the guest user device has been attached and $r$ denotes the flow rate. We use $D$ to represent the set of flow demands that have arrived in the system.

**Algorithm 1** Gateway and Path Assignment

```
Inputs: $G = (N, L), D$
SORT($D$)
for each $d \in D$ do
    $search \leftarrow$ true
    $M \leftarrow \{N \setminus n_u\}$      // $M$ : candidate set of routers
    while ($search$ AND $M \neq \emptyset$)
        $g \leftarrow argmax_{i \in M} C(n_i)$
        for each $p \in P_{ug}$
            if $C(p) < r$ then
                $P_{ug} \leftarrow P_{ug} \setminus p$
            end if
        end for
        if $P_{ug} = \emptyset$ then
            $M \leftarrow M \setminus n_g$
        else
            gateway $\leftarrow n_g$        // gateway assignment
            $path \leftarrow$ SP($P_{ug}$)     // path assignment
            $search \leftarrow$ false
        end if
    end while
end for
```

The algorithm assigns a gateway and the shortest path to each flow demand from the set $D$. The algorithm is executed whenever there is insufficient shared bandwidth in the local home network. Initially the flow demands are sorted based on the flow rate in decreasing order. For each flow demand, the algorithm selects the home router with the highest access link bandwidth (i.e., $C(n_i)$). Subsequently, the algorithm identifies the set of paths, $P_{ug}$, between the local home router ($n_u$) and the assigned gateway ($n_g$) that satisfy the flow rate requirement (Equation 1). In case there is no such path, the algorithm performs another iteration for the selection of the gateway, excluding the previously selected home router. Otherwise, the shortest among these paths is being identified based on the number of hops (*SP* function in the algorithm). This eventually designates the path for the traffic redirection through the WMN.

This algorithm carries out the gateway assignment based on the available shared bandwidth of the home routers in the crowd-shared network. Although this can pool resources from all home networks where sharing is permitted achieving efficient utilization of the shared bandwidth, guest user traffic may have to traverse multiple hops in the WMN. This can lead to latency inflation, delay variation, reordering ([16],[17]) and in general, unpredictable performance. To mitigate this, we further use a variant of this algorithm by introducing a threshold in the number of hops between the local home router and the assigned gateway. For example, adjusting the threshold to 1 restricts the search space to the next-hop routers with sufficient shared bandwidth.

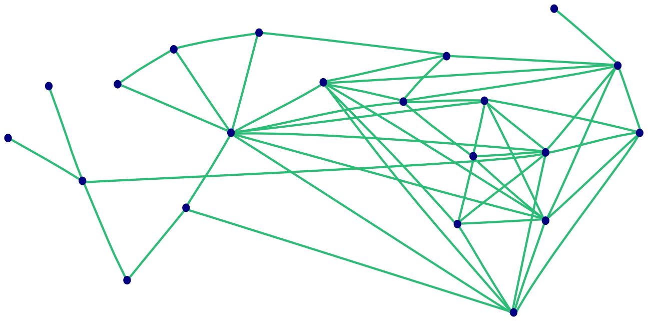

Fig. 4. Simulation WMN topology.

## IV. Evaluation

In this section, we quantify the benefits of using a crowd-shared WMN for public Internet access. In Section IV-A, we present the simulation environment, the modeling of router on/off periods, and the metrics used to evaluate the WMN efficiency. Section IV-B provides a comparative study between a crowd-shared network, such as PAWS, and a WMN, based on our simulation results.

### A. Simulation Environment

We developed a simulator that models the behavior of data flows that are generated by guest users in a crowd-shared WMN. We use the TFA wireless mesh topology [4] which consists of 21 nodes (Fig. 4). Each node in this figure represents a home router with shared Internet access bandwidth, whereas each edge is a wireless mesh link. Guest user flows arrive randomly at different home routers. Each generated flow has a rate and lifetime sampled out of a uniform and an exponential distribution, respectively. The guest user flows arrive to the network according to a Poisson process.

We model the availability of the home access routers using an on-off Markov chain. On and off times are exponentially distributed with mean values $\mu_{on=106}$ minutes and $\mu_{off} = 555$ minutes. We parameterize the exponential distributions using datasets from the PAWS deployment in UK. Fig. 5 shows the number of active routers along 60 hours of simulation. We can see that out of 21 routers less than 12 routers are simultaneously available.

Guest user flows are granted Internet access either through local routers or by redirection to remote routers based on the algorithm presented in Section III. Flows which cannot be accommodated are rejected. We evaluate the benefit of WMN in the context of crowd-shared Internet access using

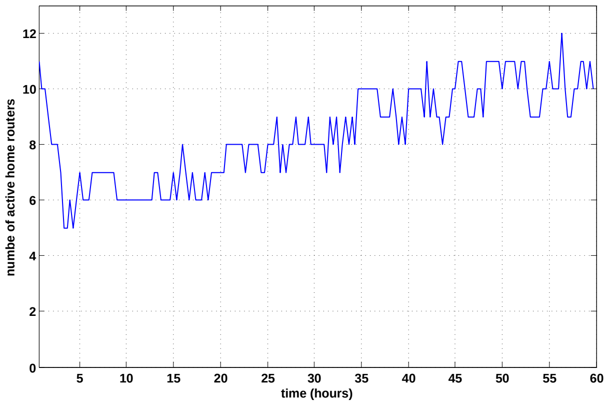

Fig. 5. Number of active home routers.

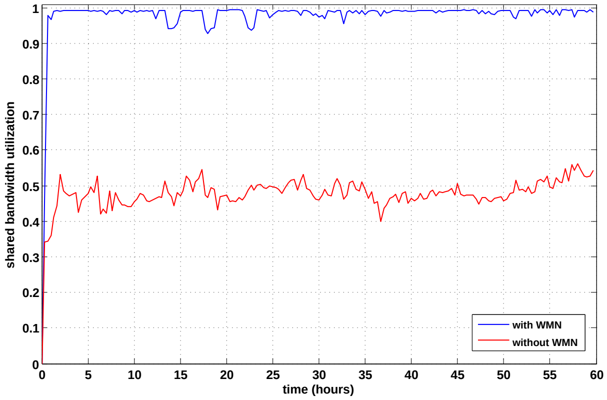

Fig. 6. Shared bandwidth utilization.

the total shared bandwidth utilization in all home routers and the accumulated serving rate $ASR(T)$ at time $T$, defined as:

$$ASR(T) = \frac{\Sigma_t^T s_t^{finished}}{(\Sigma_t^T s_t^{finished} + \Sigma_t^T s_t^{rejected})}, \qquad (2)$$

where $s_t^{finished}$ denotes the total sizes of the flows at time $t$ which are accepted (i.e., assigned Internet access) and successfully served without disruption due to router unavailability. $s_t^{rejected}$ denotes the total sizes of the flows at time $t$ which are rejected (i.e., not assigned Internet access).

Each simulation run comprises 60 hours at a time discretization of 20 minutes. For each scenario, we perform 20 simulation runs. We assume a homogeneous setting where each router has 16 Mbps ADSL downlink and set the shared bandwidth per router to 8 Mbps (i.e., in the periods that each router is active). We run our simulation with mesh link capacity of 200, 54 and 10 Mbps; this does not have any impact on the simulation results, since the bottleneck is the Internet access links. Note that our wireless link model does not consider the impact of interference and distance on the offered link bandwidth.

## B. Simulation Results

Initially, we measure the shared bandwidth utilization and serving rate with an arrival rate of 50 flows per minute over a 60-hour period. Fig. 6 illustrates a low utilization of the shared bandwidth without a WMN during the whole period, although there is high demand for Internet access by guest users attached to the various home networks. In contrast, a WMN allows to capitalize the unused capacity and accommodate a larger volume of guest user traffic. More precisely, according to Fig. 6 guest user traffic redirection through the WMN results in the full utilization of the shared bandwidth. Furthermore, crowd-shared WMNs can accommodate substantially larger volume of guest user traffic, as depicted in Fig. 7. This stems from the high utilization of the shared bandwidth.

We further measure the shared bandwidth utilization with a wide range of guest user traffic demands. In this respect, Fig. 8 illustrates the shared bandwidth utilization with diverse flow arrival rates, ranging from 10 to 100 flows per minute. This simulation result corroborates the efficiency of the WMN for various traffic loads, as the shared bandwidth utilization always remains very high. Fig. 8 shows poor bandwidth utilization without a WMN, especially with low guest user traffic demand. In this particular case, the limitation of one point of Internet access for each guest user leads to wasting most of the shared bandwidth. Essentially, our simulation results show the significant benefit that a WMN can bring into crowd-shared

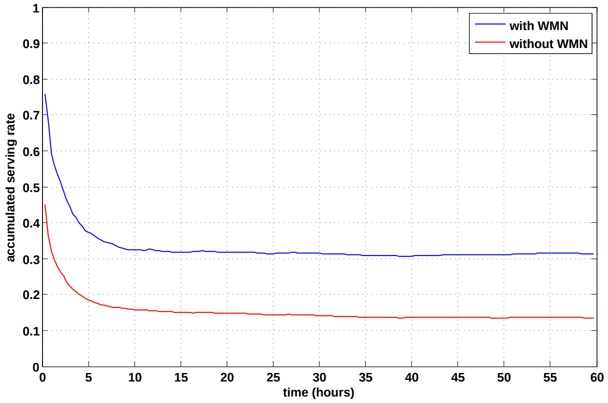

Fig. 7. Accumulated serving rate.

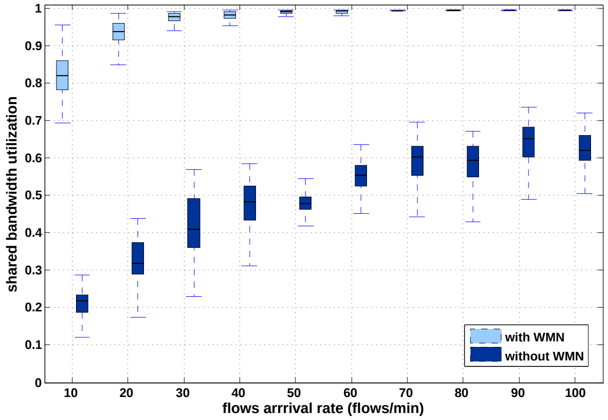

Fig. 8. Shared bandwidth utilization vs. flow arrival rate. Boxes show interquartile ranges and the median, while whiskers show the $5^{th}$ and $95^{th}$ percentiles.

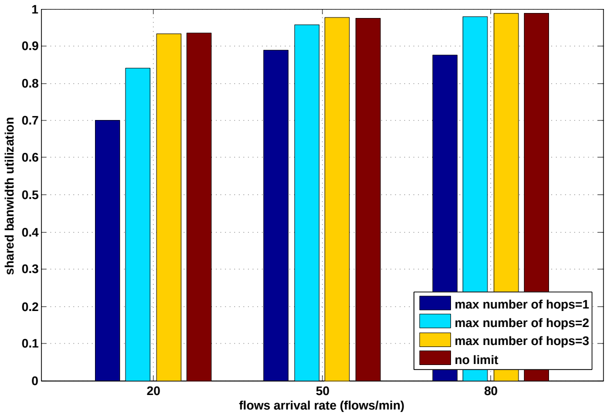

Fig. 9. Shared bandwidth utilization for diverse hop-count threshold values vs. flow arrival rate.

networks, by effectively pooling resources across all home networks.

According to our simulations, redirected guest user traffic traverses up to 4 hops in the WMN. We further use a variant of our gateway and path assignment algorithm (briefly discussed in Section III) which introduces a threshold in the number of hops between the home router and the assigned gateway. Fig. 9 illustrates the shared bandwidth utilization for diverse hop-count threshold values ranging from 1 to 3, and without any limit in the hop count. Restricting the number of hops has a noticeable impact on bandwidth utilization, especially for a single hop, since Internet access is permitted only through one of the next-hop home routers. In essence, there is a trade-off between coverage extension (and thus effective bandwidth utilization) and latency. This may become more critical in large WMNs, where multi-hop wireless links can inflate latency.

## V. Related Work

Recent work has been leveraging on SDN for the management of wireless mesh networks. In [9], Hasan et al. discuss the benefits of using SDN to decouple the operation and management of rural wireless networks from the physical infrastructure. Our work is in the same direction, but we focus on the management of crowd-shared WMN in residential areas and investigate techniques for the efficient utilization of the shared bandwidth. Authors in [6] present an OpenFlow-based architecture for WMNs with emphasis on mobility management. Dimogerontakis et al. [7] propose a SDN extension to WMN community networks for L2 experimentation. Other applications of SDN to the wireless network domain include OpenRoads [21] and Odin [19], which provide solutions for mobility management using OpenFlow. These techniques are complimentary to our work and can be integrated into our SDN control plane to facilitate mobility management in crowd-shared WMNs.

## VI. Conclusions and Future Work

In this paper, we quantified the benefits of extending the coverage of any crowd-shared network (e.g., PAWS) by connecting the home routers as a mesh. A crowd-shared WMN can mitigate the fundamental problem of any crowd-shared network, i.e., the presence of a single point of access for each guest user. In a crowd-shared WMN, the path redundancy towards the Internet backhaul can be exploited to achieve better resource utilization, especially during periods of limited Internet access sharing. According to our trace-driven simulations, a crowd-shared WMN can offer much better utilization of the shared bandwidth, providing free Internet access to a larger number of users.

SDN can greatly facilitate the configuration and management of crowd-shared WMN by third-party virtual network operators. Our SDWMN control plane provides the means for the coordination of guest user traffic redirections through the WMN. As part of future work, we will evaluate experimentally the efficiency of the SDWMN control plane in the CONFINE community network testbed.

## VII. Acknowledgments

This work was supported by the EU CONFINE project (EU FP7 288535).

## References

[1] CONFINE Project, http://confine-project.eu/.
[2] FON, http://corp.fon.com/.
[3] Freifunk, http://freifunk.net/.
[4] TFA Rice, http://tfa.rice.edu/
[5] B. Braem et al., A Case for Research with and on Community Networks, ACM SIGCOMM CCR, 43(3), July 2013.
[6] P. Dely, A. Kassler, and N. Bayer, Openflow for wireless mesh networks, IEEE ICCCN, Maui, USA, August 2011.
[7] E. Dimogerontakis, I. Vilata, and L. Navarro, Software Defined Networking for community network testbeds, IEEE WiMob, October 2013.
[8] J. Fildes, UN reveals global disparity in broadband access, BBC, http://www.bbc.co.uk/news/technology-11162656, 2010.
[9] S. Hasan, Y. Ben-David, C. Scott, E. Brewer, and S. Shenker, Enhancing rural connectivity with software defined networks, ACM DEV, New York, NY, USA, 2013.
[10] E. Kohler, R. Morris, B. Chen, J. Jahnotti, and M. F. Kasshoek, The Click Modular Router, ACM Transaction on Computer Systems, 18(3), 2000.
[11] N. McKeown et al., OpenFlow: Enabling Innovation in Campus Networks, ACM SIGCOMM CCR, 38(2), April 2008.
[12] B. Pfaff, J. Pettit, T. Koponen, K. Amidon, M. Casado, and S. Shenker, Extending Networking into the Virtualization Layer, ACM HotNets, New York, USA, October 2009.
[13] A. Sathiaseelan et al., Public Access WiFi Service (PAWS), Digital Economy All Hands Meeting, Aberdeen, October 2012.
[14] A. Sathiaseelan, C. Rotsos, C.S. Sriram, D. Trossen, P. Papadimitriou, J. Crowcroft, Virtual Public Networks, IEEE EWSDN, Berlin, Germany, October 2013.
[15] A. Sathiaseelan and J. Crowcroft, Internet on the Move: Challenges and Solutions, ACM SIGCOMM CCR, 43(1), January 2013.
[16] A. Sathiaseelan and T. Radzik, Reorder Detecting TCP, IEEE HSNMC, Estoril, Portugal, July 2003.
[17] A. Sathiaseelan and T. Radzik, Reorder Notifying TCP (RN-TCP) with Explicit Packet Drop Notification (EPDN), IJCS, Wiley, 19(6), September 2005.
[18] A. Sathiaseelan and J. Crowcroft, LCD-Net: lowest cost denominator networking, ACM SIGCOMM CCR, 43(2), April 2013.
[19] L. Suresh et al., Towards Programmable Enterprise WLANs with Odin, ACM SIGCOMM HotSDN, Toronto, Canada, August 2011.
[20] L. Townsend, A. Sathiaseelan, G. Fairhurst, C. Wallace, Enhanced broadband access as a solution to the social and economic problems of the rural digital divide, Journal of Local Economy, 28(6), 2013.
[21] K. Yap et al., OpenRoads: Empowering Research in Mobile Networks, ACM SIGCOMM, Barcelona, Spain, August 2009.