Intelligent Wifi Access Points for Diverse User needs: QoS Slicing in SDN Controlled APs

SRI HARSH AMUR, University of Twente, The Netherlands

Non-discriminatory access points pose a challenge in today's Internet that consists of diverse traffic patterns with different network expectations. Quality of Service (QoS) Slicing is a technology that is being actively researched to solve this issue. It involves creating slices inside an Access Point (AP) based on the network traffic QoS requirements. With the twofold increase of Internet traffic, both in terms of magnitude and variety, QoS slicing used in queuing and scheduling in AP is a promising way to meet user requirements. This paper implements a platform to dynamically manage slices in an SDN-controlled network with the help of traffic rules. This system improves the network performance by prioritising slices based on their application's needs. It makes use of the DSCP value of an IP header to dynamically create slices and implements traffic rules to keep the number of slices to a minimum. This paper enables better bandwidth utilization by catering to the needs of real-time, delay-intolerant slices before the delay-tolerant slices.

Additional Key Words and Phrases: QoS Slicing, Wireless Resource Management, Software Defined Networking, 5G-Empower.

1 INTRODUCTION

The number of devices connected to the internet has been rapidly increasing since the genesis of the Internet. In the past few years, IoT devices have been in mainstream use and have been exponentially surging every year. This has led to the Internet handling various kinds of traffic patterns where each pattern has a different network expectation. For example, in a smart office, many IoT devices would be communicating with each other at various times and these devices would send small packets which are delay-tolerant and don't require a lot of bandwidth. At the same time in a conference room, the conferencing app would be operating on huge packets bi-directionally requiring high bandwidth and low latency. The 2 devices significantly differ from each other in terms of network requirements, yet an Access Point (AP) or a router treats packets coming from both networks in the same manner.

Quality of Service (QoS) is a mechanism which ensures the application's requirements are being met under different networking conditions. In wired networks, the DSCP value, located in the IP header, is used to represent the application type. There are 64 possible categories using DSCP. These categories can be used by a router or AP to prioritise packets. In WiFi, QoS is implemented with the help of access categories (AC) which operate in layer 2 [12]. There are 4 AC: Voice, Video, Best Effort and Background. However, there is a need to implement a QoS mechanism in AP to make use of these categories and prioritize the packets accordingly.

QoS Slicing is a technology in which the packets are placed in different slices within an AP based on the application's requirements. These slices can be prioritised in a manner to achieve the best end-user experience. High-priority slices would get more time to transmit their packets. An SDN-controlled AP with QoS slicing can help in prioritising real-time traffic over delay-tolerant traffic and enhance the overall traffic performance. It could also perform user prioritisation in cases where different users are using a similar application.

WiFi networks are used to implement such a system over LTE, 5G or other technologies since WiFi extends to a longer range with high bandwidth. In addition, it already has a large user base, therefore such a system can be easily adopted. This paper demonstrates the benefit of utilising dynamic QoS Slicing with the help of the 5G-Empower controller [3]. Following are the contributions of this research thesis:

- *Flow Monitoring*: An application is made in the 5G-Empower controller which monitors the packet's DSCP and other fields.
- *Dynamic Slice Management*: The platform is using packet statistics to dynamically create or modify slices in the WTP.
- *Traffic Rules*: To enhance the performance of the WTP, the number of active slices is kept to a minimum with the help of traffic rules. Traffic rules combine multiple slices into a single slice when the traffic load is less.

Overall, using these 3 techniques, the platform provides great control over the bandwidth each slice gets which results in a more satisfactory QoS performance.

2 OTHER WORKS

Slicing is a key architecture of 5G and has been actively researched for the past few years. However, most of the research has been about network slicing or Infrastructure Shared Slicing [10] where the concept of network virtualization is applied over a common network infrastructure and different slices are given to different tenants. Most of the research lies within the topic of slice scheduling [9] where a scheduling algorithm is used to transmit resources between different slices. Although this is an important topic which still faces many challenges, this paper remains limited to slice management based on the classification of network traffic's QoS requirements.

Authors in [10] have used network slicing and Proportional Time Deficit Round Robin (PT-DRR) and shown that by using such a mechanism they were able to give each user a fair amount of bandwidth based on the slice priority. However, the authors focused on using network slicing instead of QoS slicing. Although our paper does not focus on AP routing algorithms, implementing QoS slicing could prove to show fair usage as well as keep the QoS requirements satisfied.

In [5], [11], the authors were able to show slicing based on QoS prioritization does not disrupt the network quality and instead allowed more efficient resources utilization. It showed that by differentiating services by QoS, the SDN network's performance improved. It reinforced how prioritising based on QoS improved the AP's performance by showing that even for uncontested APs, giving preference

TScIT 37, July 8, 2022, Enschede, The Netherlands
© 2022 University of Twente, Faculty of Electrical Engineering, Mathematics and Computer Science.
Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

to high priority slices over delay-insensitive slices, the network's QoS was delivered.

In [4], the authors made use of 5G-empower and Traffic Rules to show (i) slices are isolated from each other and do not influence each other. (ii) The resource distribution among the slices is consistent with the goodput of the networking device. (iii) The number of slices in the networking device does not affect communication. The authors also showed that modifying slices during run-time did not cause any miscommunication with other slices nor were any packets dropped. However, this research was done on LTE devices whereas this paper tries to create slices in WTP or AP. Although this research showed the dynamic nature of slicing, the authors had to manually add the slice for their experimentation. The aim of this paper is to have a system which assesses the network conditions and dynamically creates the slices without the consultation of a human.

3 BACKGROUND INFORMATION

To gain a better understanding of the problem and the solution used to tackle it, the knowledge of a few concepts is explained.

3.1 Software-Defined Networking (SDN)

SDN brings centralized, intelligent routing into a network. It essentially separates the control pane and the data pane. The data pane is a term used to describe the process of physically forwarding a packet from a networking device. This functionality remains inside the WiFI router or AP. However, control pane is a term used to describe the process of deciding where and how the packet will be forwarded. This functionality will be embedded into the controller. The controller sits as a centralized node in the network, having a top-down view of the entire network in order to orchestrate it. A note to make is that packets do not travel from the router or AP to the controller for the controller to make a decision. Rather statistics about a bunch of packets are sent to the controller in short intervals. In this paper, 5G-Empower is used to make a controller and to make an LVAP agent (Light Virtual Access Point) which is installed in the AP. The term Wireless Terminal Point (WTP) is used for an AP with an LVAP agent inside it.

3.2 DSCP

Differentiated Service (DiffServ) is a mechanism used to classify packets to enable better QoS for the end-users. In the header of an IP packet, the DiffServ field uses 8 bits out of which the first 6-bits are used for Differentiated Service Code Point (DSCP) and the last 2 bits are used for Explicit Congestion Notification (ECN). The DSCP is used to classify packets based on the QoS requirements and is used by various networking devices for prioritization. Using the DSCP value, a networking device might drop or delay a packet while facing congestion. RFC 2474 [2] defined 4 classes for the DSCP values.

3.2.1 Default Forwarding (DF). DF or Best Effort (BE) represents the default behaviour of the networking device. The numerical value of DF is 0. A packet with 0 in its DSCP field is given no special treatment.

3.2.2 Expedited Forwarding (EF). EF is used by packets that need low latency and are delay-sensitive. The decimal value of EF is 46. Usually, real-time services use this since they need low latency and low packet loss. Video streaming, conference calls and audio calls are examples of services which should have the EF value in their DSCP field. Although it depends on how the router or AP is configured, these packets would be given a high priority.

3.2.3 Assured Forwarding (AF). AF is used for packets that need delivery assurance. But the packets using AF can be dropped if the network is facing congestion or if the arrival rate has exceeded the network capacity. To address how the AF packets should be dropped, the 6 bits used to represent the AF are split into 2 categories. The first set of 3 bits describes the sub-class and the value ranges from 1 to 4. The second set of 3 bits is used to represent the drop-preferences and the value ranges from 1 to 3. The higher the sub-class value, the higher the priority. Similarly, the higher the drop preference value, the more likely that packet will be dropped. Therefore, in total there are 12 different values for the AF class. A packet with DSCP value AF23 is prioritized over a packet with the value AF12. This is because although the drop preference of the first packet is higher but since its sub-class value is higher, it takes a higher priority.

3.2.4 Class Selector (CS). Originally, the DiffServ field used to be called the Type of Service (ToS) field. The first 3 bits were used to let the network know the packet priority and were called IP-Preference. To ensure DSCP is backwards compatible with IP-Preference, CS was used as a class. There are 7 CS values ranging from 1 to 7. CS1 is usually used for irrelevant packets and is commonly known as scavenger packets. CS2 to CS4 are used by ordinary packets. CS5 is used for high-priority packets and CS6, and CS7 are used for network management packets.

3.2.5 Voice Admit. In RFC 5865 [8], a new class was defined called Voice Admit which is similar to EF and its DSCP value is 44.

The values that are not defined in the RFC are usually treated as Best Effort. A point to note is that although these are the recommended values and classes, it is up to network administrators on how they want to treat each DSCP value.

However, in practice, it has been observed that the majority of all IP packets transmitted through the internet use the Best Effort (0) value in their DSCP field. This goes to show that the DSCP of IP packets is not accurate and a better classification system needs to be used to assign a more accurate DSCP.

3.3 QoS Slicing

[6] A flow is defined as a data stream which transmits from a source to a destination and can be identified using its source port number, source IP address, destination port number, destination IP address and transport layer protocol. An AP can handle multiple flows at the same time. A slice is when one or more flows are grouped together. A flow can only belong to one slice but the end-user can generate multiple flows and therefore be part of multiple slices concurrently.

Network slicing or Infrastructure Sharing Slicing (ISS) is another type of slicing that has been introduced as a key architecture of 5G wherein virtual networks can be created on top of a common infrastructure. These network slices or virtual networks can be used

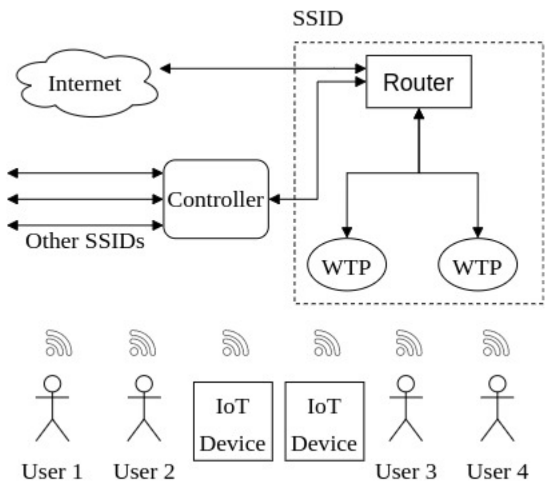

Fig. 1. Network Set-up.

by tenants to host their own network which is isolated and serves their client's requirements. However, In this paper, slicing would be done within an AP to serve different flows. The criteria to group together flow into a slice is if the flows share similar QoS requirements. Hence the name QoS slicing. To know the QoS of the flow, flow properties like packet sizes, packet arrival rates, port numbers and IP addresses can be used to classify the traffic. A slice in 5G-Empower is identified by the SSID (WiFi network identifier) and DSCP. An AP would place packets into different slices depending upon the DSCP value of the packet. These slices would be given different priorities to fulfil the function of QoS assurance. The scheduling algorithm of the AP would correspondingly get packets from these slices following the priority.

4 METHODOLOGY

The set-up consists of a controller and a few WTPs like in Figure 1. The controller would receive statistics from all WTPs in the network about the DSCP of the packets. Using this, it gives instructions to all the WTPs.

4.1 Slice Grouping

In this paper, the scheduling used for transmitting packets is called Airtime Deficit Round Robin(ADRR). With ADRR, each slice is given a quantum value which is used as a unit of measure to transmit packets. Every packet that is transmitted subtracts a certain amount of deficit from the quantum and the packets from a slice are transmitted until the quantum value is less than what the packet requires. For example, a slice with a 2000 quantum value gets twice as much air time to transmit its packets when compared to a slice with a quantum value of 1000.

However, Having a different slice for each DSCP leads to very small quantum values resulting in reduced air time per slice. Instead, grouping together slices of different DSCPs into one slice is a better alternative since this enables each slice to transmit for a significant amount of time. Therefore, a mechanism is needed which can combine multiple slices into a single appropriate slice using an SDN framework. Using such a slice grouping mechanism, many packets with different DSCP values would be inserted into a few slices and controlling the air time for each slice would be easier. The groupings used in this paper are shown in Table 2.

Table 1. Priority of Slice Groups

Priority	Class Name	Quantum Proportion
1	Network Management	4
2	Real-Time	3
3	Streaming	2
4	Broadcasting	1.5
5	Best Effort	1
6	Scavenger	0.5

In this paper, we follow the RFC standard and keep the priority as stated. This priority is shown in Table 1.

As can be seen, we have 6 slice groups. These can be further broken into their respective DSCP slices dynamically if the traffic demands. The number of groups is also more than what WiFi Access Control offers [12], which offers only 4 groups: Voice, Video, Best Effort and Background. Having more groups allows the controller to render finer control over the complete network.

The quantum proportion is used to assign quantum values to each slice. The total quantum value assumed to be possible is 10,000. Whenever a slice is created or re-evaluated, based on the group slice and quantum proportion, it is given a quantum. For example, If there are 3 slices to be created, one for BE, one for Streaming and one for Real-Time, then the BE slice would be given a quantum of 1666, the Streaming slice a quantum of 3333 and finally the Real-Time slice a quantum of 5000.

4.2 Slice Management

Based on the Slice grouping, the controller checks if those slices already exist in the WTPs. In case they do not exist, the controller would send instructions to WTPs. Slice removal is not implemented in the system yet since additional slices that do not have any packets do not get any airtime. Therefore, having additional slices in the WTP does not ruin the efficiency significantly. However, in the future, having such a system would enable better control over the entire network.

4.3 Traffic Rules

The Slice groupings happen in the controller and are sent to the WTP. The WTP is running the LVAP agent [3] using the Click Modular Router (CMR) [7] according to which "CMR is a software architecture for building flexible and configurable routers. A Click router is assembled from packet processing modules called elements." The Traffic Rule element is responsible for ensuring group DSCP are applied.

A traffic rule consists of a match and an action. A match is a dictionary of source IP address, source port, destination IP address, destination port, protocol and DSCP code value. The action is a

Table 2. Slice Groupings

DSCP Value	Class	Standard Use-Case	Group DSCP	Group DSCP Class
0	BE	Best Effort	0	Best Effort
8	CS1	Scavenger	8	Scavenger
16	CS2	Network Control	0	Best Effort
24	CS3	Broadcasting	24	Broadcasting
32	CS4	Streaming	32	Streaming
40	CS5	High Priority	24	Broadcasting
48	CS6	Network Management	48	Network Management
56	CS7	Network Management	48	Network Management
10	AF11	High Throughput Data	0	Best Effort
12	AF12	High Throughput Data	0	Best Effort
14	AF13	High Throughput Data	0	Best Effort
18	AF21	Low Latency Data	0	Best Effort
20	AF22	Low Latency Data	0	Best Effort
22	AF23	Low Latency Data	0	Best Effort
26	AF31	Multimedia Broadcasting	24	Broadcasting
28	AF32	Multimedia Broadcasting	24	Broadcasting
30	AF33	Multimedia Broadcasting	24	Broadcasting
34	AF41	Multimedia Conferencing	32	Streaming
36	AF42	Multimedia Conferencing	32	Streaming
38	AF43	Multimedia Conferencing	32	Streaming
44	Voice Admit	Voice Calls	46	Real-Time
46	EF	Real-Time Interaction	46	Real-Time

DSCP code value. The Traffic Rule element tries to match each packet with existing rules. If a match is found then the action is performed where it would change the DSCP value in order to place different packets into the same slice.

4.4 DSCP Statistics

The DSCP statistics is an element in the WTP's LVAP agent. It is responsible for collecting the packet information. This information is sent to the controller as a response to every request made by the controller.

The statistics include the number of packets per DSCP. This is used to judge if packets of a particular DSCP are increasing or decreasing and whether to create a separate slice for a DSCP. Another statistic is the average size of a packet per DSCP. This can be used to check if the packet size of a DSCP is very large. Since large packets take more time to transmit, it could be better to create a new slice for a DSCP with large packets in some situations. This ensures that packets of other DSCPs do not get bottle-necked. Lastly, packet information is also sent consisting of the packet's source and destination IP address and port number, protocol and DSCP. This statistic is not used in the current system, however, it can be used in the future to make even more fine-grained traffic rules for a particular flow. For example, if a flow has to be given priority over another flow of the same slice, then using these packet features, the flow can be switched to another slice with a higher priority. These statistics can also be used to estimate the slice throughput requirements. Additionally, knowing the throughput requirements helps

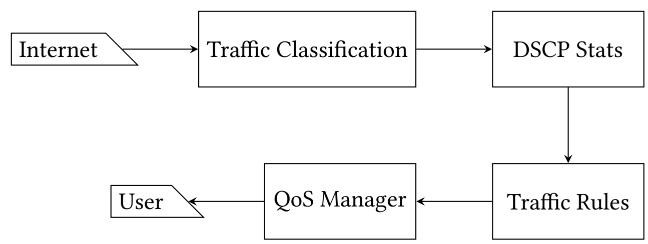

Fig. 2. Packet Flow inside the WTP.

in developing a quantum assignment algorithm. This algorithm can be used to calculate the quantum value of each slice dynamically.

4.5 Packet Flow and Architecture

Figure 2. shows the packet flow inside the WTP. The packet arriving from the internet first needs to be assigned a more appropriate DSCP. This is due to the fact that the majority of the packets arrive with a DSCP of 0 [1]. An element placed at the beginning of the flow can inspect a packet and assign it a more appropriate DSCP. The current system uses the packet features to assign it a DSCP. However, a machine learning model can be used instead in the future for traffic classification and accurate DSCP assignment. Following that, the DSCP Stats element stores the packet information. After which the traffic rule element checks if the packet matches any rules and applies the rule subsequently. Each slice is identified by the pair *(SSID, DSCP)*. Therefore, by observing the DSCP of the packet, the QoS manager finally places the packet into the corresponding slice.

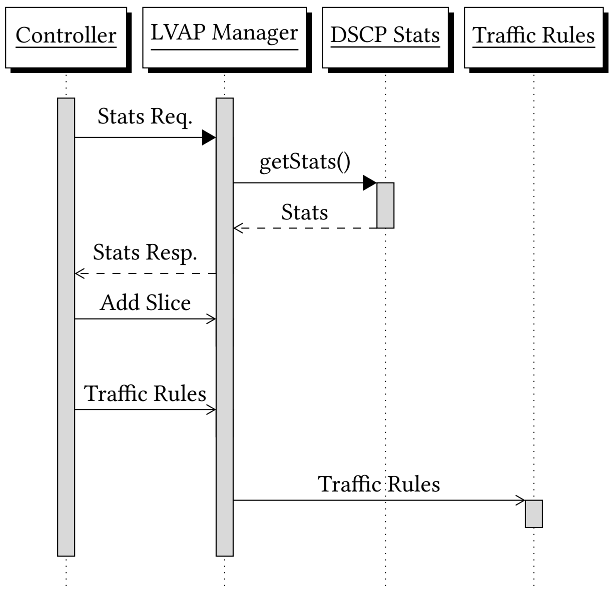

Fig. 3. Interaction between Controller and WTP.

Fig. 3 shows the interaction between the controller and the WTP. It starts with the controller requesting the statistics. The LVAP manager collects the statistics from the DSCP Stats element and sends them back to the controller. The controller gets these statistics from all active WTPs and forms an idea of the whole network. It then analysis these statistics to form slice groups and traffic rules. It then either adds slices if they do not already exist or updates the existing slices. In addition, it also sends the traffic rules to the LVAP manager so that packets are put into the slice groups instead of making a slice for each DSCP.

5 TESTING FRAMEWORK

To show that QoS Slicing enables prioritization in the WTP during congestion, we organised 2 different scenarios. Both are framed in such a way that the WTP faces congestion after a certain time and needs to prioritise which packets to transmit and which ones to drop.

5.1 Scenario 1

The network has one WTP and one controller. At the start, i.e 0 seconds in Fig. 4 6 Best Effort TCP Flows are transmitted to the WTP. They are all kept in a single slice and no prioritization takes place as of yet. After 30 seconds, 2 new flows are added. These flows have a DSCP of 32 and 38 respectively. Since these indicate streaming and conferencing applications, they need to be given a higher priority. The controller should make a new slice for DSCP 32 with a higher quantum value than BE and should also make a traffic rule to change DSCP 38 into 32 so that both flows utilise the same slice. After the 1-minute mark, 2 more flows are added. These flows have a DSCP of 44 and 46. They indicate voice calls and real-time interaction and should be prioritized over the other 8 flows. Therefore, the controller makes a new slice for DSCP 46

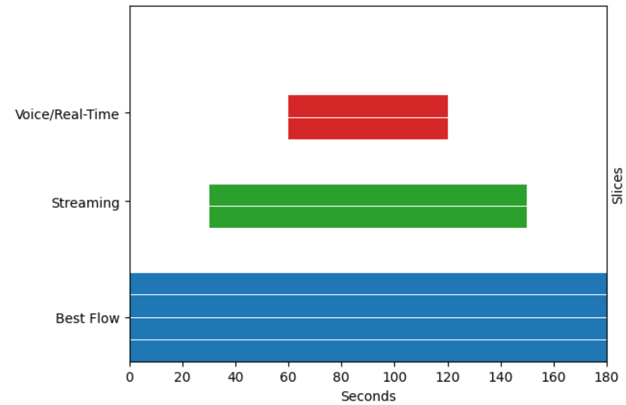

Fig. 4. Experiment 1 Set Up

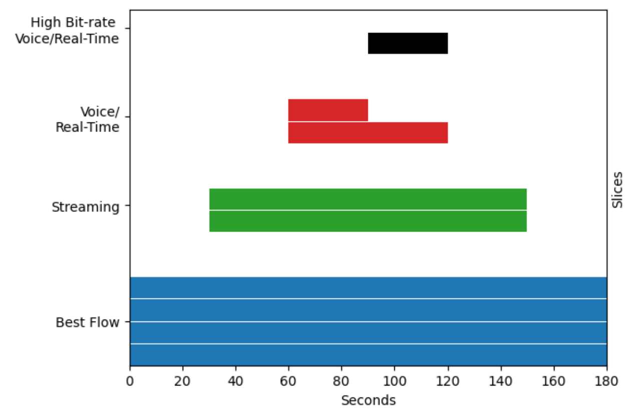

Fig. 5. Experiment 2 Set Up

with a higher quantum than both the other slices. In addition, a new traffic rule should be made to change DSCP 44 into 46. After the 2-minutes mark, the latest 2 flows are now removed (DSCP 44 and 46). The controller needs to re-organise the slices with different quantum values, after which the configuration goes back to that of the 0 to 1-minute mark.

The experiment is conducted to test if the controller can notice the congestion and flow conditions to dynamically add slices and create traffic rules.

5.2 Scenario 2

The second scenario is made to test the dynamic nature of slice creation and slice groupings. The network has one WTP and one controller. At the start in Fig. 5, 4 Best Effort TCP flows (Flow A) are transmitted to the WTP. They should all have the same priority since they would be kept in one slice. After 30 seconds, 2 new flows are added with the DSCPs 32 and 38 which are streaming and conferencing packets (Flow B). After the 1-minute mark, 2 new flows with DSCP 44 and 46 (Flow C) are added. So far the AP should have 3 slices, one for Best Effort, one for Streaming/Conferencing and one more for Real-Time flows. At the 90th second, the flow

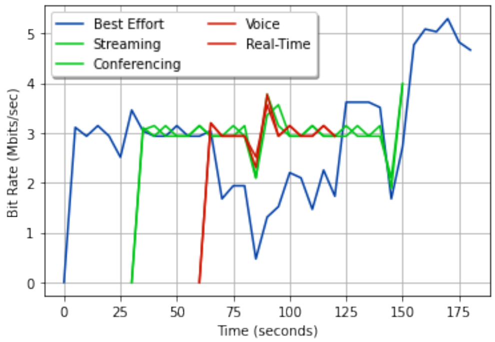

Fig. 6. Experiment 1 with Slicing

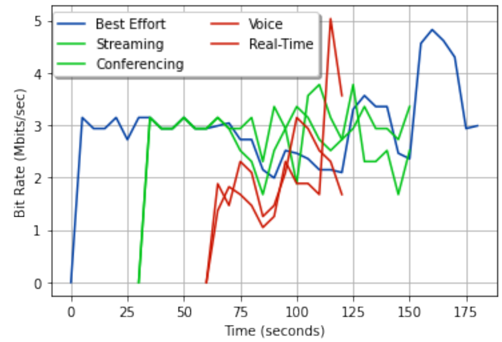

Fig. 7. Experiment 1 without Slicing

with DSCP 44 increases its data rate. The controller should notice that slice 46 is filling up due to the increase in packets of DSCP 44 and should create a separate slice for 44. In addition, it should also remove the previously existing rule by overriding it. After the 2-minutes mark, Flow C should stop and the controller should be able to re-organise the slices with suitable quantum values. Finally, at the 150th second, Flow B should stop as well.

This experiment tests if the controller is able to ensure when a flow with a certain DSCP value increases beyond a threshold, the new slice for that flow is created to cater to those packets. This shows that packets with higher priority are not dropped due to an increase in the flow of other packets (up to a point).

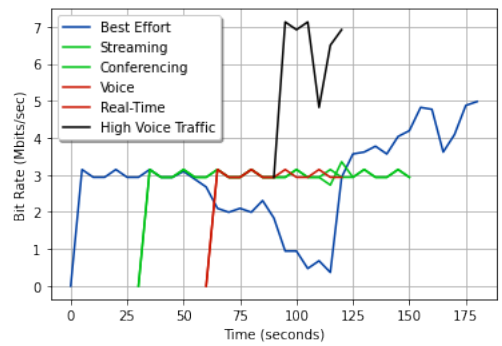

Fig. 8. Experiment 2 with Slicing

6 RESULTS

The bit rate expectation of each flow was set to 3 Mbits/second. Each experiment was done twice, once with slicing enabled and once without slicing. This was done to check the benefits slicing provides.

6.1 Slice Grouping and Prioritization

In Fig. 6 Flow A starts and stabilises at around 3 Mbits/second. Flow B also stabilises at 3 Mbits/second. However, once Flow C starts, the network congestion begins. The AP checks which flow has the highest priority and caters to its bit rate first. Therefore, Flow C gets 3 Mbits/second. Similarly, since Flow B has second priority, it also gets 3 Mbits/second. Finally, the remaining bit rate is provided to Flow A. After Flow C is terminated, due to the available bit rate, Flow A shoots up. But soon the AP reorganises the slices to increase Flow B's rate.

Fig. 7 shows the result of not using slicing with such a set-up. Since the AP does not know how to prioritize the flows, it randomly assigns a bit rate to each flow thus causing Flow C to get an average of 2 Mbits/second instead of 3. Slicing gives the AP the ability to manage resources in a much more efficient and intentional way. Based on the slice priority, the AP can ensure the bandwidth is distributed proportionately.

6.2 Dynamic Slice Creation

In Fig. 8 Flow A and B start one after the other and stabilise at 3 Mbits/second. At the one-minute mark, Flow C joins the networks.

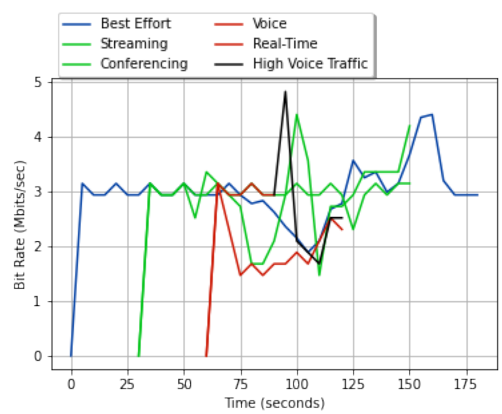

Fig. 9. Experiment 2 without Slicing

This leads to congestion and the AP reduces Flow A's bandwidth to cater to Flow B and C. At the 90th second, one of the flows in Flow C increases its bit rate to 7 Mbits/second. Since this is quite high, the controller makes a new slice for that flow and gives it the same priority as Flow C. This further decreases the bandwidth of Flow A and ensures the high priority slice's network expectations are met.

In Fig. 9 the same network flow without slicing is shown. It firstly does not give priority to Flow C which leads to a delay in packet arrival. This would disrupt the voice quality of the end-users. Secondly, when the Flow demanded a higher bandwidth of 7 Mbits/second, the AP was able to provide a peak rate close to 5 Mbits/second and then dropped it back to 2 Mbits/second immediately.

7 FUTURE WORK

In the future, the following aspects of the platform can be improved:

- *Traffic Classification*: A traffic classification element can be placed inside the WTP to accurately assign an appropriate DSCP. Since the platform depends on the DSCP of the packet, a good model could identify different QoS requirements and accordingly create more slices.
- *Dynamic Slice Grouping*: The current platform has fixed slice groupings and can future break a slice group into its respective individual slices. However, due to the unpredictable nature of traffic patterns, a more dynamic grouping strategy could be beneficial in a real-world scenario.
- *Dynamic Quantum Assignment*: The quantum values for each slice group are currently fixed. However, a machine learning model which dynamically assigns quantum values to a slice depending on other factors could optimize the platform more efficiently and can make better use of the bandwidth.

8 CONCLUSIONS

QoS Slicing has shown the potential to manage a network with diverse traffic patterns. This paper implements a platform using 5G-Empower to dynamically manage QoS slices in an SDN-controlled network. It uses the DSCP value in the IP header to create slice categories and assigns these slices different quantum values to schedule packets according to a priority. It uses slice grouping to group together slices which have similar QoS in order to keep the number of slices low which in turn leads to significant quantum values for each slice. To ensure slice groupings are applied in the WTPs, it makes use of Traffic Rules to change the packet's features.

The results show that the bandwidth is distributed proportionally according to the slice priority. High-priority slices are given enough bandwidth to transmit their packets before other slices. For example, low-latency, delay-intolerant slices get the bandwidth to send their packet for a longer time since any disturbance caused otherwise degrades the experience of the end-user. The paper has also shown that if a flow requires higher bandwidth than other flows in the same slice, it is given its own slice. This way other flows' bandwidth is not consumed by this flow and the new slice's bandwidth can be determined dynamically by taking into consideration the network congestion.

REFERENCES

[1] Davide Aureli, Antonio Cianfrani, Alessio Diamanti, José Manuel Sanchez Vilchez, and Stefano Secci. 2020. Going Beyond DiffServ in IP Traffic Classification. In *NOMS 2020 - 2020 IEEE/IFIP Network Operations and Management Symposium.* 1–6. https://doi.org/10.1109/NOMS47738.2020.9110430

[2] Fred Baker, David L. Black, Kathleen Nichols, and Steven L. Blake. 1998. Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers. RFC 2474. https://doi.org/10.17487/RFC2474

[3] Estefanía Coronado, Shah Nawaz Khan, and Roberto Riggio. 2019. 5G-EmPOWER: A Software-Defined Networking Platform for 5G Radio Access Networks. *IEEE Transactions on Network and Service Management* 16, 2 (2019), 715–728. https://doi.org/10.1109/TNSM.2019.2908675

[4] Estefania Coronado and Roberto Riggio. 2019. Flow-Based Network Slicing: Mapping the Future Mobile Radio Access Networks. In *2019 28th International Conference on Computer Communication and Networks (ICCCN).* 1–9. https://doi.org/10.1109/ICCCN.2019.8847068

[5] Kevin Ibarra-Lancheros, Gustavo Puerto-Leguizamón, and Carlos Suárez. 2018. Quality of service evaluation based on network slicing for software-defined 5G systems. *TecnoLógicas* 21 (09 2018), 27–41. https://doi.org/10.22430/22565337.1066

[6] Vinod Joseph and Brett Chapman. 2009. *Introduction to Cisco's Quality-of-Service Architecture for IP Networks.* 23–70. https://doi.org/10.1016/B978-0-12-374461-6.00002-1

[7] Eddie Kohler, Robert Morris, Benjie Chen, John Jannotti, and M. Frans Kaashoek. 2000. The Click Modular Router. *ACM Trans. Comput. Syst.* 18, 3 (aug 2000), 263–297. https://doi.org/10.1145/354871.354874

[8] James Polk, Fred Baker, and Martin Dolly. 2010. A Differentiated Services Code Point (DSCP) for Capacity-Admitted Traffic. RFC 5865. https://doi.org/10.17487/RFC5865

[9] Matías Richart, Javier Baliosian, Joan Serrat, and Juan-Luis Gorricho. 2016. Resource Slicing in Virtual Wireless Networks: A Survey. *IEEE Transactions on Network and Service Management* 13, 3 (2016), 462–476. https://doi.org/10.1109/TNSM.2016.2597295

[10] Matías Richart, Javier Baliosian, Joan Serrat, and Juan-Luis Gorricho. 2020. Resource Allocation and Management Techniques for Network Slicing in WiFi Networks. In *NOMS 2020 - 2020 IEEE/IFIP Network Operations and Management Symposium.* 1–6. https://doi.org/10.1109/NOMS47738.2020.9110407

[11] Iana Siomina and Stefan Wanstedt. 2008. The impact of QoS support on the end user satisfaction in LTE networks with mixed traffic. In *2008 IEEE 19th International Symposium on Personal, Indoor and Mobile Radio Communications.* 1–5. https://doi.org/10.1109/PIMRC.2008.4699672

[12] Tim Szigeti, Jerome Henry, and Fred Baker. 2018. Mapping Diffserv to IEEE 802.11. RFC 8325. https://doi.org/10.17487/RFC8325