

A DiffServ transport network to bring 3G access to villages in the Amazon forest

Javier Simó-Reigadas



PHOTO: Radio station in Rio Napo , Peru. ONGAWA.



CASE STUDIES A DiffServ transport network to bring 3G access to villages in the Amazon forest

EDITED BY

Global Dimension in Engineering Education

COORDINATED BY

Agustí Pérez-Foguet, Enric Velo, Pol Arranz, Ricard Giné and Boris Lazzarini (*Universitat Politècnica de Catalunya*)
Manuel Sierra (*Universidad Politécnica de Madrid*)
Alejandra Boni and Jordi Peris (*Universitat Politècnica de València*)
Guido Zolezzi and Gabriella Trombino (*Università degli Studi di Trento*)
Rhoda Trimmingham (*Loughborough University*)
Valentín Villarroel (*ONGAWA*)
Neil Nobles and Meadhbh Bolger (*Practical Action*)
Francesco Mongera (*Training Center for International Cooperation*)
Katie Cresswell-Maynard (*Engineering Without Border UK*)

DL B 10735-2015
ISBN 978-84-606-7546-4

This publication is distributed under an Attribution- Noncommercial- Share Alike License for Creative Commons



Citation: Simó-Reigadas, J. (2015) 'A Diff-Serv transport network to bring 3G access to villages in the Amazon forest', in *Case studies for developing globally responsible engineers*, GDEE (eds.), Global Dimension in Engineering Education, Barcelona. Available from: <http://gdee.eu/index.php/resources.html>

Disclaimer: This document has been produced with the financial assistance of the European Union. The contents of this document are the sole responsibility of the authors and can under no circumstances be regarded as reflecting the position of the European Union.

A DIFFSERV TRANSPORT NETWORK TO BRING 3G ACCESS TO VILLAGES IN THE AMAZON FOREST

Javier Simó-Reigadas, Universidad Rey Juan Carlos.

INDEX

1. INTRODUCTION	3
1.1. DISCIPLINES COVERED	3
1.2. LEARNING OUTCOMES	4
1.3. ACTIVITIES	4
2. DESCRIPTION OF THE CONTEXT	5
2.1. URBAN AND RURAL AREAS IN THE DEVELOPING WORLD	5
2.2. THE PERUVIAN CASE	5
2.3. THE TUCAN3G PROJECT	8
2.4. SCENARIO OF THE CASE STUDY	9
3. CLASS ACTIVITY	10
3.1. PRESENTATION OF THE CONTEXT OF THIS CASE STUDY	10
3.2. QUALITY OF SERVICE (QoS) IN IP NETWORKS	11
3.3. THE DIFFSERV ARCHITECTURE	14
3.4. QUEUING DISCIPLINES	19
3.5. SOLUTION AND EVALUATION CRITERIA	22
4. HOMEWORK ACTIVITY	23

1. INTRODUCTION

Since their conception, Internet Protocol (IP) networks have been generally used under the best-effort paradigm. Several Quality of Service-aware (QoS-aware) IP network architectures have been proposed in the last two decades by the Internet Engineering Task Force (IETF) and other actors, DiffServ being the most successful. Modern network technologies covering layers 1 and 2 such as Asynchronous Transfer Mode (ATM), Ethernet, WiFi or WiMAX have included mechanisms for traffic differentiation and some support for QoS. However, it is still uncommon to find a scenario where a QoS-aware network architecture, such as DiffServ, is adopted and implementation is simple enough to be fully understood by undergraduate students.

This case study helps the student to understand what problems arise in an IP network when several terminals that need QoS guarantees share an IP infrastructure, and how the DiffServ architecture helps to solve the problem. Moreover, a practical implementation of queuing disciplines in the network nodes is proposed. The student will learn how a simple embedded computer with Linux may become a DiffServ router, and how this can be applied to propose an appropriate communications solution for rural deployments in developing regions. This case study is based in the TUCAN3G project, a FP7 research project that aims to propose a low-cost solution for 3G coverage in rural areas of developing countries (TUCAN3G, 2013). It comprises a theoretical session about the fundamentals, followed by a problem resolution activity in which the theoretical concepts must be applied to a well-defined scenario.

This case study is not intended to provide the student with a deep understanding of DiffServ and queuing disciplines. Students must have a good understanding on IP networks prior to this case study, and it may be useful if they have already studied DiffServ and queuing disciplines from a theoretical perspective. The case study helps students to integrate all related information and associate the functionality of these tools with the needs identified in a real life scenario.

1.1. DISCIPLINES COVERED

This case study covers the implementation of a QoS-aware IP network. The case study proposes a clear example in which a 'common' IP network is not appropriate, and the use of the DiffServ architecture converts the network in an appropriate solution. The case also shows how to implement it, and demonstrates a case in which the local support of DiffServ is valuable, no matter what happens beyond. The clues for implementing

DiffServ nodes are simple and permit any student to go further on his/her own by building low-cost DiffServ routers.

1.2. LEARNING OUTCOMES

- The student will learn how DiffServ contributes (qualitatively and quantitatively) to make an IP network QoS-aware.
- The student will learn how to identify QoS limitations in an IP network, and to propose appropriate actions to improve the overall performance.
- The student will learn how to apply well-known queuing disciplines to implement the Per-Hop Behaviour (PHB) desired for a particular case in a DiffServ domain.
- The student will learn how QoS-aware IP networks based on WiFi and WiMAX may become a key solution for the deployment of low-cost communications infrastructures in rural areas in developing regions.

1.3. ACTIVITIES

The following learning activities are proposed:

1. Theoretical session: 2 hour class on QoS (Quality of Service), the DiffServ architecture and the implementation of a basic DiffServ solution. The case study presented is a means to understand the detailed performance problems to be tackled in an IP network in which certain traffic classes require QoS guarantees. Prior to this session, the student will have read basic materials on QoS, DiffServ (referenced below) and queuing disciplines, or have received basic training on these subjects. It is also advisable for students to read the context of this case study in advance.
2. Problem resolution activity: the students are organised in working groups and presented with a communications problem based on the same context explained in the theoretical session. Several basic communications elements and techniques are proposed, and each group discusses the best way to solve the problem, as well as define all details for network implementation. The outcome is a technical report produced by each group.

2. DESCRIPTION OF THE CONTEXT

2.1. URBAN AND RURAL AREAS IN THE DEVELOPING WORLD

Urban areas are growing all around the world. Moreover, the fact that more than 60% of the population are based in urban areas attracts economic investment. On the other hand, rural areas contain only 40% of the population, but cover a much larger area, yet contain so providing this population with services is comparatively difficult and expensive (Klein, 2007). Most rural areas lack basic infrastructure and quite a few are very isolated from a geographic viewpoint.

In the particular case of Latin America and Caribbean, the domestic gap between urban and rural areas is large, which suggests a handicap for the integral development of those living in rural areas of the region. This is demonstrated by differences between the average and Adjusted by Inequality HDI (Human Development Index) in the region, which were 0.74 and 0,55 in 2013 respectively (UNDP, 2013). According to the Economic Commission for Latin America and the Caribbean (ECLAC), the region has over 126 million people living in rural areas, 50.2% of whom live below the poverty line (ECLAC, 2012).

The development opportunities for those millions of people in Latin America and the Caribbean are seriously constrained by the negative impact that the lack of infrastructure has on education, local production, health and so forth. Paradoxically, in many cases there are financial resources available at both the national and the local levels for investment in infrastructure, often derived from the activity of extractive industries. The problem is determining what infrastructure national policies and rural communities consider worth investing money in for the promotion of local development in rural communities, and how funds can be efficiently used.

2.2. THE PERUVIAN CASE

Peru is the third largest country in Latin America at 1,285,216 km². The population density is very low at 23.7 inhabitants/km², due to the large areas of sparsely populated jungle and mountain ranges. The census in 2007 mentioned that the population of Peru was 27,409,200, with 75.92% living in urban areas and 24.08% in rural areas. For more detail of the population distribution across communities of different sizes, see Table 1.

Table 1: Population distribution in rural and urban areas in Peru.

Size of community (no. inhabitants)	POPULATION (2007)			
	Total population	Urban	Rural	Percentage
<100	1,571,362	0.01%	5.72%	5.73%
>=100 & <300	2,977,959	0.28%	10.58%	10.86%
>=300 & <1000	2,780,110	3.04%	7.10%	10.14%
>=1000	20,079,769	72.59%	0.67%	73.26%
Total	27,409,200	75.92%	24.08%	100%

Source: TUCAN3G project, deliverable (TUCAN3G-D21, 2013), based on INEI – 2007 Census.

Regarding access to Information and Communication Technology (ICT), Table 2 shows people's access to different types of services depending on where they live. It is apparent that those living in the mountains or jungle have less access to communications services, especially access to data networks. The lack of access to communications services in rural areas such as these is also demonstrated by data presented in Table 3. Comparing the data in Tables 2 and 3 with Table 1, it can be seen that most communities with under 300 inhabitants are in those regions and lack of any communications infrastructure.

Table 2. Households with access to ICT and monthly ICT-related expenses.

Households with Access to ICT	Region			
	Coast * (%)	Mountains (%)	Jungle (%)	Metropolitan Lima (%)
Telephone (fixed)	31.11	10.97	14.39	55.69
Cellular	80.84	65.97	63.87	85.63
Internet	14.81	6.91	4.41	32.67
Average monthly household expenses (€)				
Telephone (fixed)	10.73	11.98	10.93	14.05
Cellular	8.50	7.19	9.61	10.20
Internet	18.72	18.15	21.06	21.17

Source: (TUCAN3G-D21, 2013), based on INEI-ENAH0, 2011
 * Metropolitan Lima is not considered in the coastline region.

Table 3: Population living in areas with telecommunications services coverage

Service	Urban	Rural
Fixed telephony	86%	0%
Mobile + Fixed wireless telephony	92%	53%
Fixed broadband access to the Internet (ADSL)	82%	0%
Mobile access to Internet 2.5G (EDGE)	92%	48%
Mobile broadband access to Internet (UMTS)	56%	3%
Cable TV	67%	0%
Satellite TV	100%	100%
Public telephony	94%	56%

The use of Information and Communication Services (ICT) services was measured by Peru's National Institute for Statistics and Information (INEI) in the census of poverty levels published on 2012. This census clearly showed that both poverty and lack of telecommunications services are high in rural areas, which consist primarily of relatively small communities.

On the other hand, telecommunications, and more specifically telephony, is seen by the rural population as a key factor for development. The ESAN business school recently conducted a study in more than 300 rural towns called "Evaluation of the results of FITEL (Telecommunications investment fund (Peru)) rural projects and baseline for the services continuation supported by OSIPTEL (Telecommunications regulatory institution of Peru)". According to this study, 93.8% of the rural population prefer to use mobile telephony service as opposed to any other telephony service.

Many more studies analyse the use of telecommunications in rural areas in more depth, but this basic data gives enough insight to understand that people living in rural areas are poorer, have much less infrastructure in general, and do not have access to 3G networks in particular (see Figure 1). On the other hand, people in these areas care about mobile coverage and associate their opportunities for human development to their access to telecommunications networks.

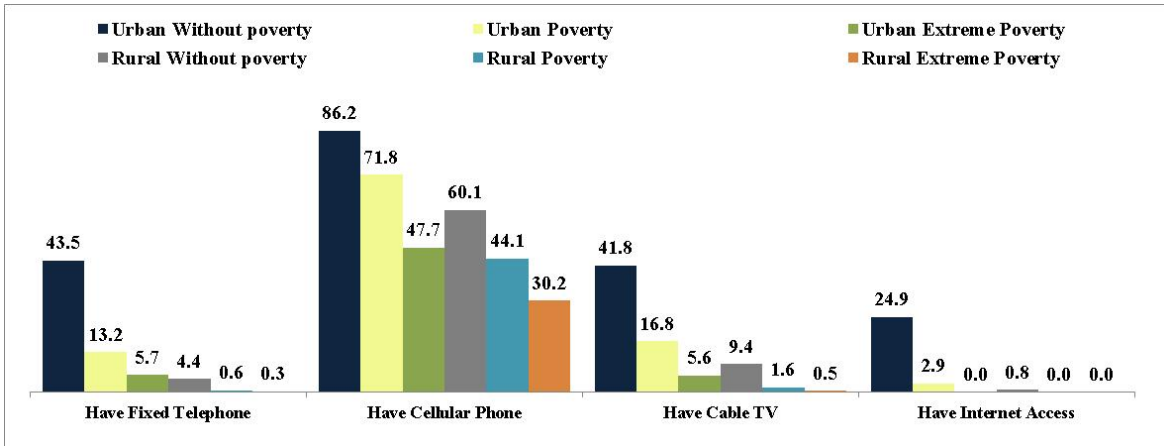


Figure 1: Population with and without coverage according to the number of inhabitants per town (at end of 2011). SOURCE: (TUCAN3G-D21, 2013)

2.3. THE TUCAN3G PROJECT

TUCAN3G is the acronym of an international research project entitled “Wireless technologies for isolated rural communities in developing countries based on cellular 3G femtocell deployments”, funded by the FP7 programme of the European Commission (TUCAN3G, 2013). TUCAN3G aims to eliminate or at least reduce the scientific, technical and economic problems that limit the deployment of current 3G or 4G systems in rural and sparsely inhabited areas. Currently these deployments result in non-profitable business to mobile operators: the well-known trade-offs between coverage and capacity imply large investments that are not associated with large revenues.

A significant cost for communications networks in rural areas is associated to the related transport network, also called “backhaul” in this case study. Together with the solutions for the access network, high-capacity low-cost solutions for the transport network are also needed. WiFi for Long Distances (WiLD) has been proposed and successfully used to connect remote regions in developing countries. WiFi systems are extremely low-cost compared to other broadband technologies, and some improvements to Medium Access Control (MAC) layer, together with the modern physical improvements included in IEEE802.11n, permit the establishment of long distance links (up to 75 Km or longer point-to-point) giving throughput rates of several Mbps. Using this technology, several multi-hop networks up to 500 kilometres have been deployed in remote areas of developing countries by members of the TUCAN3G consortium, including the Napo network mentioned below. Nowadays these networks are being employed for health and education purposes but can also be used as backhaul for providing cellular services in remote localities.

WiMAX (Worldwide Interoperability for Microwave Access) has also emerged as a solution for providing broadband communications with Quality of Service (QoS) in remote areas of both developed and developing countries. WiMAX (in both licensed and non-licensed bands) will be also considered in this project for the backhaul network, working as the only backhaul technology or creating a heterogeneous network with WiLD or the traditional Very Small Aperture Terminal (VSAT) solution.

The project aims at evaluating the viability of backhauling 3G femtocells¹ with WiLD-WiMAX-VSAT networks, from both technical and economical points of view. While femtocells are initially conceived for indoor use, interest has been recently growing for outdoor coverage due to the ease of deployment. In TUCAN3G, its use in remote rural communities is proposed as an appropriate way to provide their inhabitants with telephony and broadband connectivity.

2.4. SCENARIO OF THE CASE STUDY

Between 2006 and 2009, a linear network was deployed to connect some rural health facilities existing along the Napo river starting at Iquitos city and finishing in Cabo Pantoja, just close to the border between Peru and Ecuador. The network was implemented by the EHAS (Hispanic-American Health link) foundation and GTR-PUCP and is known as the Rural Telemedicine Network of the Napo River. Through this network the users are able to access to the Internet, telephony, data transfer and video conferencing. These services have improved the quality of health care for rural residents.

The Napo network is a chain of WiFi links, with lengths between 25 and 50 kilometres. There are 13 rural health facilities being served and 5 more relay stations. Figure 2 shows a schema of the whole network.

Although the Napo network is only used for telemedicine, a key value of this network that can be shared are the supporting infrastructures. These are towers with heights between 60 and 92 meters that ensure the LOS (line of sight) in every hop. There is no reason why these towers might not be used as the supporting infrastructures for a parallel network,

¹ For simplicity we are adopting “femtocells” in the sequel, in the understanding that it refers to small cell-related technologies based on either picocells or femtocells. Femtocells are served by femto base stations, also called Home Node B (HNB in UMTS or HeNB in LTE).

and TUCAN3G aims to exploit this possibility. The idea is to deploy a transport network that serves as backhaul for several rural 3G femtocells between Negro Urco and Santa Clotilde (see in the Figure 2). The same towers already supporting the telemedicine network are going to be used for both the access nodes and the transport network's nodes. This network is going to be deployed for demonstration purposes and is the one focused on for this case study.

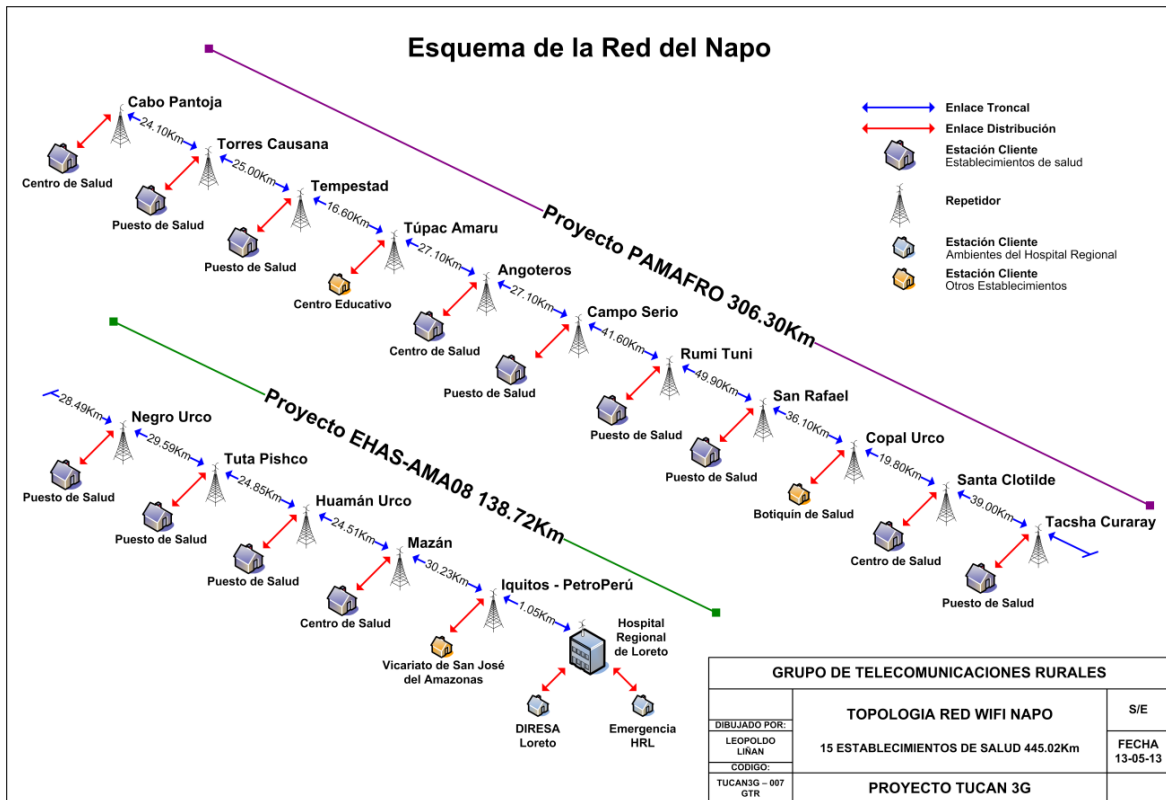


Figure 2: Napo network scheme. Note that the network is represented in two segments for graphical reason but it is a single linear network (Tacsha Curaray is connected to Negro Urco).

3. CLASS ACTIVITY

The initial class activity foreseen for this case study is divided in four parts of approximately 20 minutes each: (1) presentation of the case study, (2) revision of QoS concepts in IP networks, (3) the DiffServ architecture, and (4) queuing disciplines. This is followed by the problem resolution activity.

3.1. PRESENTATION OF THE CONTEXT OF THIS CASE STUDY

The lecturer should briefly introduce the objectives of the session and explain that a case study activity is going to be used to better understand the concepts of QoS and the DiffServ architecture. It is important that the students realise that this context must be

understood first because several examples used along the rest of the class will be based on it. The students must also know that the classroom session will be followed by a group homework in which they must apply these concepts and propose a valid solution.

The lecturer should then introduce the context of the case study and use the slides provided for introducing the scenario and the technologies used in it. The context part of this case study should be available to the students prior to the class activity, so that they can read it in advance and make this part of the class activity much more fluent and short.

3.2. QUALITY OF SERVICE (QoS) IN IP NETWORKS

The Internet was conceived in the last century as a data communication network with a best-effort philosophy, which means that the network did not assume any engagement with the users in terms of quality of the services offered. The Internet has always aimed to do just “its best effort”. On the other hand, telecommunications operators used to have large infrastructures for real-time services such as telephony in which circuit switching ensured certain service parameters. In this sense, once a circuit was established between users in the public switched telephone network the quality was ensured – except in case of exceptional events – because communications resources were dedicated for the duration of that communication.

However, in the 90's a few real-time services started to be offered in the Internet, and this trend became increasingly important in the first years of the 21st century due to their flexibility and cost effectiveness of these. However, the unpredictable nature of the Internet in terms of offered end-to-end capacity, delay, packet-loss, etc. showed that the “best-effort” behavior was clearly invalid for modern communications services such as IP telephony, video-conferencing, etc., which used to have the circuit-switching networks as the quality assurance. The concept “Quality of Service” was then defined as the objective value ensured by the network according to certain parameters. This concept is closely related with “Quality of Experience” (QoE), which is a more subjective concept related to the users' perception of the quality of services.

The International Telecommunications Union (ITU-T (Telecommunication Standardization Sector) defines the QoS as the “totality of characteristics of a telecommunications service that bear on its ability to satisfy stated and implied needs of the user of the service” [E-800]. However, the QoE is more commonly understood as a set of user-driven performance requirements. The ITU-T establishes in the Recommendation [G-1010] that the key parameters to consider for QoS are the following:

- Delay is defined as the time taken to establish a particular service from the initial user request, and also as the time to receive specific information once the service is established. Delay has a very direct impact on user satisfaction depending on the application, and includes delays in the terminal, network, and any servers.
- Delay variation is generally included as a performance parameter since it is very important at the transport layer in packetised data systems due to the inherent variability in arrival times of individual packets. However, services that are highly intolerant of delay variation will usually take steps to remove (or at least significantly reduce) the delay variation by means of buffering, effectively eliminating delay variation as perceived at the user level (although at the expense of adding additional fixed delay).
- Information loss has a very direct effect on the quality of the information finally presented to the user, whether it be voice, image, video or data. In this context, information loss is not limited to the effects of bit errors or packet loss during transmission, but also includes the effects of any degradation introduced by media coding for more efficient transmission (e.g. the use of low bit-rate speech codecs for voice).

Another parameter that is commonly considered is the throughput, understood as the rate in bits per second of data successfully sent and received. The throughput is actually a relative performance indicator, because it can only be tagged as good or bad when compared with the expected value for a given service. For example, a voice codec generating 64 kbps and producing a throughput of the same value is very good, while a video service expecting 384 kbps and obtaining only 350 kbps is very bad (although the value is higher than in the first example). When the throughput is lower than the traffic injected by the transmitter, this may increase the delay in the short-term if there are queues that can store packets temporarily, but in the long-term it will increase the information loss (because queues are finite and, once they are full, packets have to be dropped).

The G-1010 ITU-T recommendation contains a detailed relation of services with a short description of their specific requirements in this sense, and it also contains several tables with quantitative thresholds for the performance indicators of the most important services.

Depending on the specific nature of a particular communication, it may have strict requirements on all these parameters, some of them, or none of them.

- Elastic traffic: traffic flows that may adapt to the existing resources and that are very tolerant to the QoS offered. A typical case is a web page. If the throughput is higher, the web page will be presented sooner by the web navigator, and the same thing can be said if the delay or the delay variation are low. A significant packet-loss can be tackled by the transport protocol that is in charge of the retransmission of lost packets, and can also be translated into a higher delay. Obviously, the user wants to get the web page as fast as possible, but the range of acceptable values for each QoS parameter is quite large.
- Soft real-time traffic: traffic flows with some restrictions to the QoS parameters. A typical case is a video-streaming system: the throughput is defined by the codecs used, and no more throughput is generated if there are more resources available. The delay is not significant, the user may wait a little while until the video is played, but once it starts it must not be interrupted. Hence, throughput and delay variation are restricted, but the delay is generally flexible. Regarding the packet-loss, it is acceptable at a given threshold without having a great impact on the Quality of Experience (QoE).
- Hard real-time traffic: traffic flows that have hard constraints on minimum throughput, maximum delay, maximum delay variation and maximum packet-loss. The most typical case in the telephony.

DISCUSSION: The lecturer should discuss with the students about the effects of an insufficient available capacity, a higher delay, a high delay variation or a higher packet-loss than required on the QoE of users for the services mentioned (web navigation, video streaming and telephony). Students should also identify other services that can be classified as elastic, soft real-time and hard real-time.

FROM THE CASE STUDY: take the case of an operator willing to deploy 3G services in the context presented above. The operator has just deployed 3G services in Iquitos, which is one of the biggest cities in the Amazon Forest. Now they want to provide services in villages along the Napo River from Mazán to Santa-Clotilde. We will look later at the details of the transport network required for that, for now, imagine that there is IP connectivity good enough in each village. The access network consists of small cells (one or two per village) powered with solar photo-voltaic systems. Each cell exchanges three kinds of traffic with a controller that resides in the operator's network:

- **Voice traffic (telephony).** Each small cell has 16-24 channels for phone calls. Each phone call in a 3G small cell consists of two IP traffic flows: the upstream and the downstream. The cell does circuit multiplexing in the uplink for the sake of efficiency, building bigger packets in the upstream because each packet contains a fragment with voice samples for each active phone call. On the other hand, the downstream is not multiplexed, each voice packet in the downlink belongs to a single flow. Although the throughput depends on the number of active phone calls, as well as on other factors, for the sake of simplicity we are going to consider that a voice call generates 20 kbps in the upstream and 60 kbps in the downstream.
- **Signaling.** The cell exchanges signaling traffic with the controller. The rest of the traffic makes no sense unless the signaling packets are successfully transmitted and received. The amount of signaling traffic is much lower than 1% of the rest of the traffic.
- **Data traffic.** All data applications are in this block. Although this set is very heterogeneous, we are considering data traffic as flexible traffic very adaptive to the available resources.

DISCUSSION: How must each type of traffic in the case study be characterised ?

3.3. THE DIFFSERV ARCHITECTURE

DiffServ was proposed in the 1990's as a means to give an IP network QoS support (Kilikki, 1999). An alternative approach called IntServ (which stands for Integrated Services) was also proposed in the same period, but due to scalability problems IntServ was not as successful as DiffServ has been. This QoS architecture sees all the traffic in the network as formed by traffic aggregates. DiffServ does not care about individual flows, and knows nothing about the state of each communications. IP packets are just classified in groups depending on the type of traffic. For example, a voice packet belongs to the same traffic class no matter which traffic flow it belongs to. The keys for this architecture to work are two:

- Each packet must receive a mark that is respected along its way through the DiffServ domain. This mark may be set by the transmitting user or by the first system in the domain, which is called edge node. A byte in the IP header previously called ToS (Type of Service) is now called DS (Differentiated Services) and its first 6 bits are called DSCP (Differentiated Services Code Point). Although the standards recommend certain DSCP

values for different traffics, the network administrator for a DiffServ domain may assign other values as long as the criteria are consistent across the domain.

- Each node in the DiffServ domain must recognise the different marks in IP packets and must behave as appropriate with each traffic class, depending on the traffic class needs and its own capacities. Each traffic class requires a different type of QoS and, consequently, each node along a path has a well-defined per-hop behavior (PHB) that corresponds to the appropriate way to deal with the different traffic classes. For example, a web packet can be delayed and even lost without excessively compromising the communication quality, whereas a phone packet can be discarded but not greatly delayed, and a signaling packet should not be discarded. Each node implements this kind of logic in its PHB. For a more detailed study of DiffServ, use the references given below.

There are basically three different types of PHB:

- Best Effort PHB (BE): all packets receive the same treatment under this PHB, and there is no specific guarantee related to performance indicators.
- Expedited Forwarding PHB (EF): this traffic is known to have a high priority and to be sensitive to delay and delay variation. Packets are given priority and packet schedulers reserve resources to ensure that this traffic experiences a low delay.
- Assured Forwarding PHBs (AF_{xy}, where x is an integer between 1 and 3 and y is an integer between 1 and 4). Traffic classes receiving this PHB have priority over BE traffic. The priority is higher or lower depending on the value of 'x'. Depending on how important is for a traffic class to suffer from packet-loss in the case of congestion, 'y' may be given a higher or lower value.

A complementary concept used in DiffServ is the SLA (Service Level Agreement), which is the QoS assured by the operator to the client. It may include a maximum throughput (peak and average values may be explicitly negotiated), and sometimes a maximum packet-loss rate, a maximum delay and a maximum delay variation.

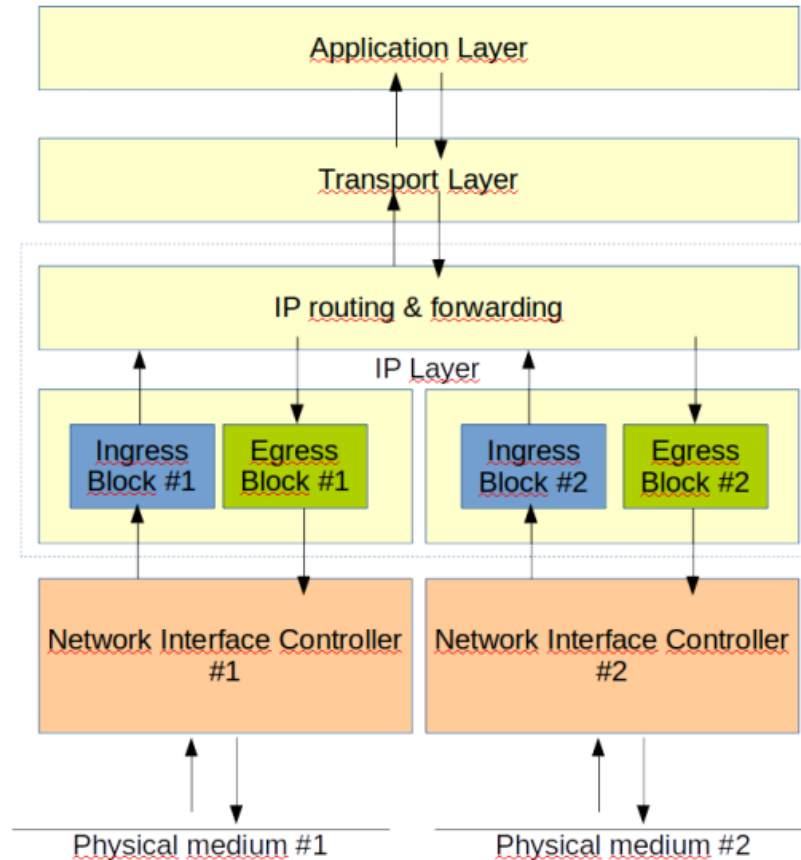


Figure 3: Flow of IP packets in an IP router.

There are several techniques that an IP router needs to use for a full implementation of its PHBs, in this case study all the techniques are not developed in-depth. Instead, the most important tools that permit the easy obtainment of acceptable behavior are considered, which is developed in a group homework activity. First, some basic concepts are introduced.

An IP router may have several NICs (Network Interface Controllers), that may be interfaces to an Ethernet network, an Asynchronous Transfer Mode (ATM) network or whatever network technology the router is connected to. Packets arriving to the system from an NIC get to an “ingress block” in which the system may apply a **policy** (i.e. certain packets may be discarded in order to enforce the incoming traffic to accomplish the Service Level Agreement, or SLA), **classification & marking** (i.e. to distinguish among traffic classes), **traffic shaping** (i.e. limiting the transmission rate in a controlled way), **scheduling** (how much bandwidth is given to each traffic class and how much priority each one receives). In general, any actions may be applied to a single packet. Each packet that is accepted in the system goes to a block in which the destination address is

examined in order to decide if the packet must be forwarded to another system or must be passed to the upper layer instead. In the former case, the routing table is used to determine the outgoing Network Interface Controller (NIC) to which the packet must be forwarded, and then the packet passes to its egress block. An egress block contains a queuing discipline in which outgoing packets are queued. There are many types of queuing disciplines, in the reference below, the student may revise the most important disciplines and understand what each one is used for.

DISCUSSION: The students should discuss where in the diagram above a system applies its PHB.

FROM THE CASE STUDY: Remote 3G cells are spread over sparsely populated areas. Each cell has very few simultaneous users. The use of a common infrastructure for the transport network lowers its cost, and the low demand of each individual small cell makes it possible to do this.

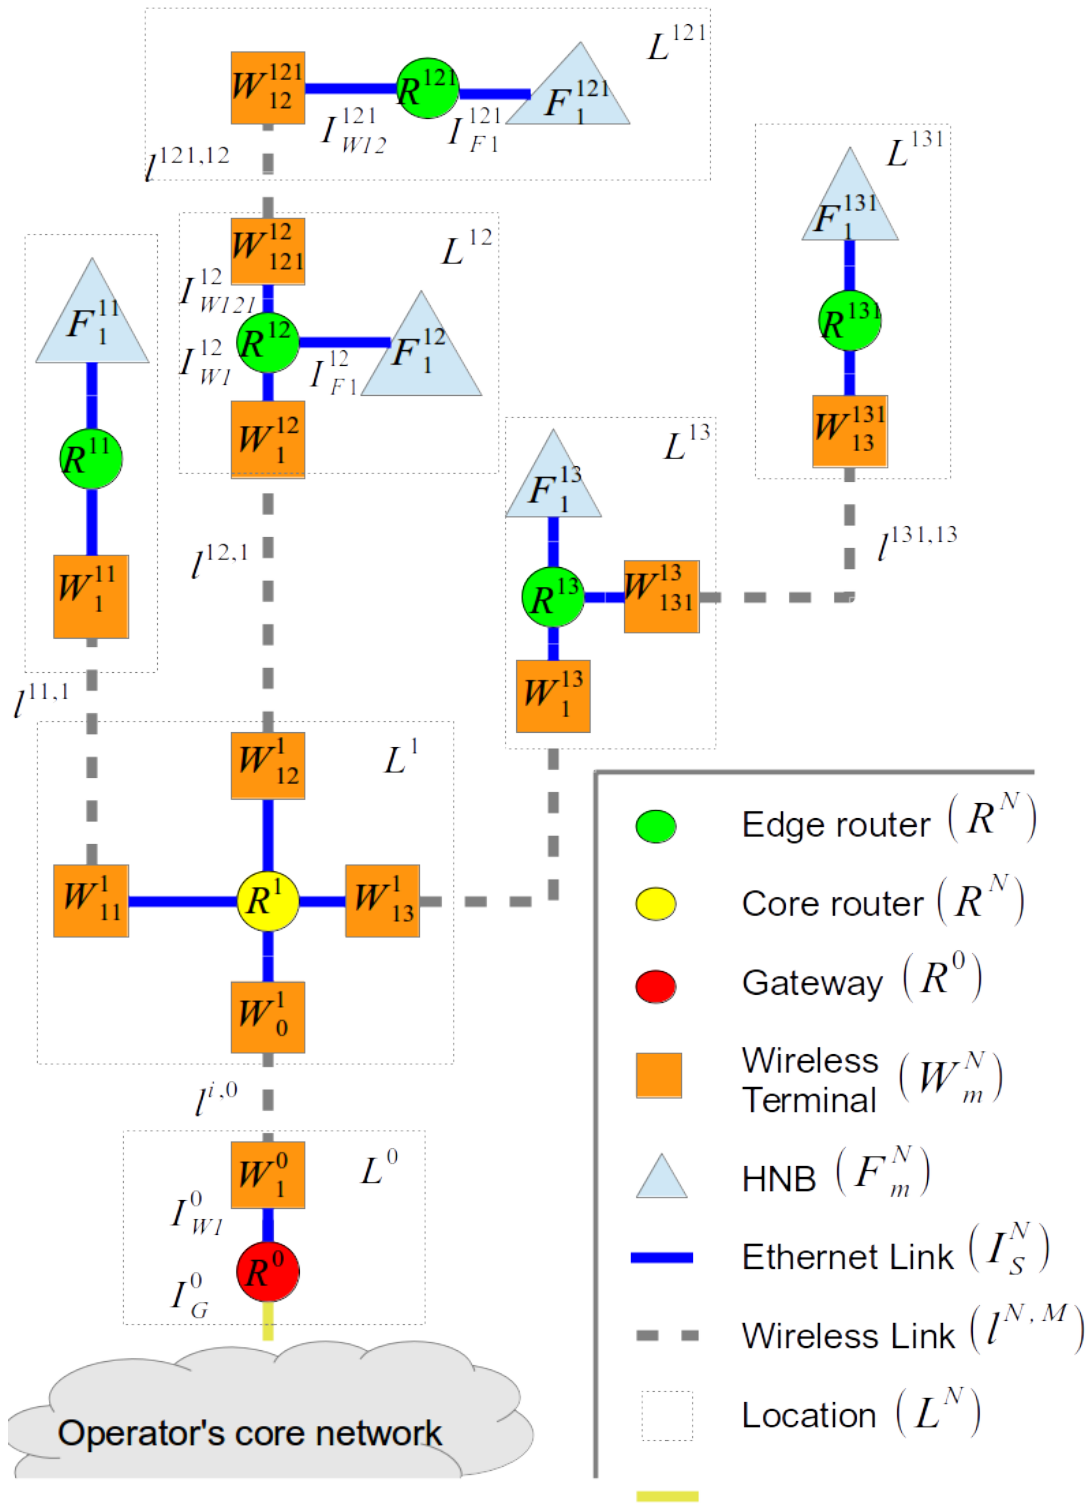


Figure 4. Transport network architecture

In Figure 4, the transport network topology that better fits the case study is a tree. The “roots” of the tree are the nodes that connect the transport network to the operator’s core network (we call it “transport network gateway”). Small cells are represented as blue triangles, and the nodes they are connected to are called “edge nodes”. Nodes are all IP

routers and are represented by circles. In this project, research activities have shown that WiFi or WiMAX links may be used to link any pair of nodes that are fifty kilometres away or even more. The only important condition is that the router has a limit for the maximum throughput that it can present to the link, in order to ensure that the link works with very low delay, negligible packet loss and stable behavior. Hence, each node must receive packets, decide what link it must forward them to, classify them, differentiate traffic classes and give them different priorities, and restrict the maximum throughput that is presented to each node. This is done with DiffServ. The small cells (called HNBs in the figure) mark all IP packets with a DSCP that indicates the traffic class each packet belongs to.

DISCUSSION: What should nodes do with the traffic? In what block -ingress, egress- should they apply these operations? Take Figure 3 and try to specify how things happen as much as possible, provided that the whole transport network is a DiffServ domain.

3.4. QUEUING DISCIPLINES

IP routers switch IP packets, which means that individual IP packets get into the router through an interface (or are originated by upper layers in that router). Figure 3 shows how packets are guided through the system. Once packets are sent to the egress block for transmission, the system must put them in a queue so that packets are not lost if the channel is busy. The problem of QoS then arises, because packets in a queue are subject to several problems:

- If a packet has several other packets in front of it in the queue, it is going to be significantly delayed. Depending on how constant the delay experienced by consecutive packets in the same flow is, this may impact on the delay, the variation of the delay, or both.
- If the number of packets arriving to the queue is higher than the number of packets that can be transmitted in average, the queue will eventually be full and packets will be dropped. A packet arrived to a full queue is dropped, no matter how important it is.

In order to combine the necessity of queuing the packets and the QoS requirements, a complex queuing discipline may be required (LARTC, 2014).

A queuing discipline (also commonly called qdisc) determines how packets are managed. The simpler qdisc is a FIFO queuing discipline that puts each new packet always at the

end of the line. A FIFO qdisc has a certain size (that can be described either in number of packets or in bytes) and packets arriving when the queue is full are dropped. Each time a packet can be effectively transmitted through the channel, the first packet in the line (usually called the 'HOL', the head-of-line) is pulled from the queue and a new position becomes available at the end.

The FIFO qdisc is the simplest. There are other qdiscs that have more than one band. In this case, there are different queues, there are filters that determine which packets go into each queue, and there is a scheduler that decides that the HOL must be taken for the next transmission in the channel.

Some qdiscs may be even more complex and have classes. A qdisc has classes if each band may in turn be another arbitrary qdisc (with one or several bands, with one or several classes).

The use of such complex qdisc may be complicated but the possibilities for obtaining a precise behavior in an IP router are also enormous. A qdisc may serve to do traffic conditioning, establish priorities between traffic classes, adjust the tradeoff packet-loss / delay, etc.

Any practical IP router supports a variety of qdiscs, and this is also the case for the Linux kernel. Moreover, the fact that Linux is the most extended open-source operating system has driven many researchers and developers to use it to implement, test and improve all IP control techniques. Hence, Linux is one of the most powerful options for practising with qdiscs and is the option considered in this case study. In the referenced document LARTC (2014), the student may find a comprehensive description of how Linux deals with IP packets, what qdiscs and complementary tools exist, how they work and how to use them for practical purposes. The students (and the lecturer) may read this document, available in several languages, in order to gain a complete perspective on qdiscs and how to use them.

The following figures are examples of some qdiscs that are very similar to those that the students use for the problem resolution activity proposed in this case study. Figure 5 illustrates a Hierarchical Token Bucket qdisc that limits the outgoing traffic to a maximum bitrate of 15,000 kbps. It gives maximum priority to voice traffic and reserves 5,800 kbps for this traffic class. Then it reserves 100 kbps for signaling traffic, and the rest is given for any other traffic flows with a minimum guarantee of 1 Mbps. The ToS byte of each packet

(which coincides with the DS byte) is read and passed in order to decide in what child qdisc the packet must be queued. Due to their regularity, voice packets are put in a FIFO qdisc. The other types of traffic are put in Stochastic Fair Queuing (SFQ) qdiscs, which have several bands and avoid blocks occurring among different traffic flows.

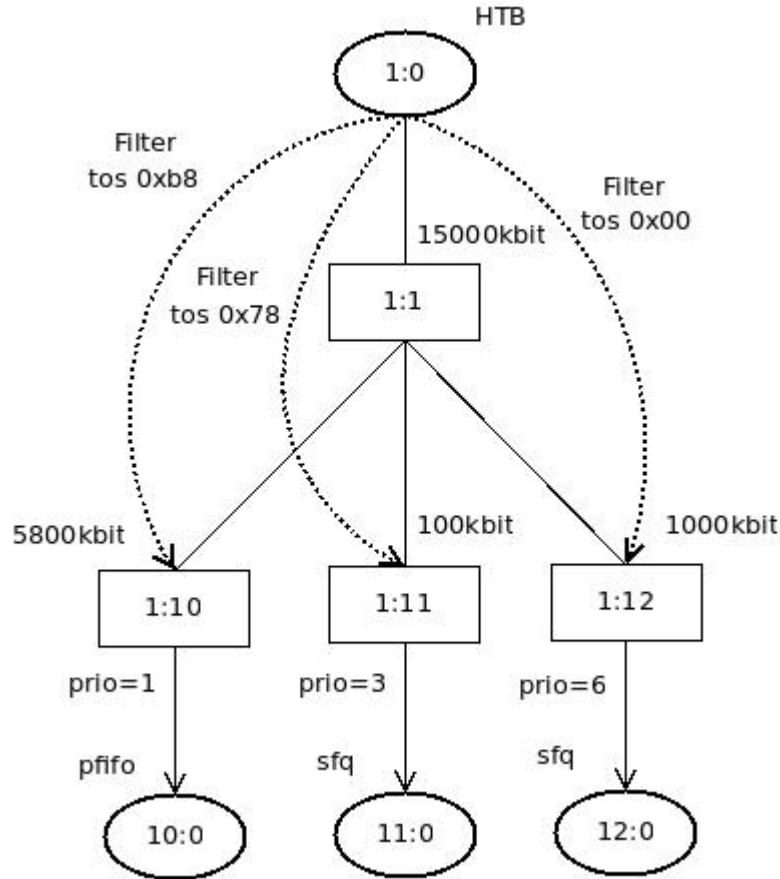


Figure 5: Diagram illustrating a HTB queuing discipline with PFIFO and SFQ child qdiscs. The value of TOS to look for is indicated in the filters, the bandwidth assigned to each band is also defined, and so are the priorities. Different configurations are expected for the different systems.

Figure 6 shows another case of complex qdisc: PRIO. It has several bands and applies absolute priorities. In this case, the highest priority is always given to voice packets, no matter how many packets are waiting in the other queues. The signaling traffic is processed, and data traffic is only sent when the other two queues are empty.

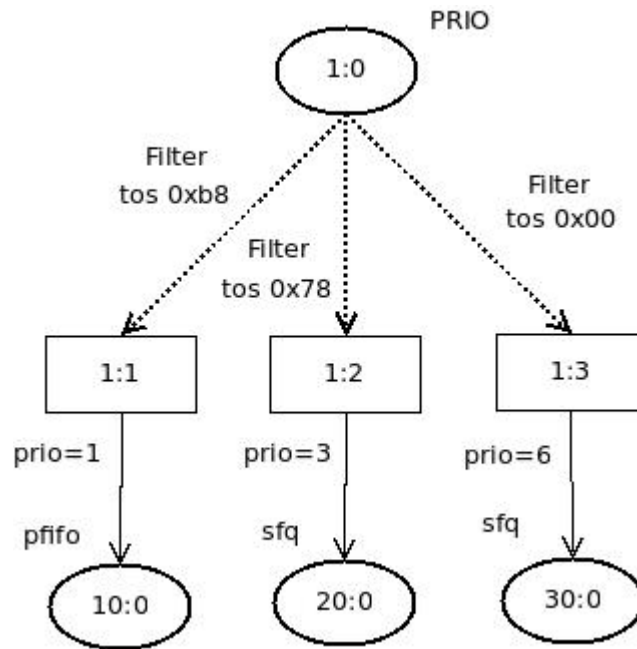


Figure 6: Similarly to the previous Figure, a PRIO alternative is represented here.

3.5. SOLUTION AND EVALUATION CRITERIA

TUCAN3G has worked intensively for demonstrating that both WiLD and WiMAX point to point links can be valid solutions for the wireless interconnection of any pair of nodes in the transport network used for backhaul. The project has also proved that:

- The IP network operating as best-effort is not a valid solution because
 - links tend to present high delays (specially in the case of WiFi, due to retransmissions and queues) or high packet-drop probabilities (in both WiFi and WiMAX when any link operates beyond the saturation point).
 - Voice traffic requires much more control on its QoS, and especially on its delay.
 - Signaling traffic requires much more control on its QoS, and specially on its packet-loss.
 - A best-effort network presents unfairness among access nodes. Those closer to the gateway obtain more resources than those that are further away.

- Several QoS-aware solutions have been examined. However, some solutions under the IP layer (such as Multi-Protocol Label Switching, or MPLS) and other IP architectures such as IntServ are not implementable in low-cost rural networks because standardisation bodies and manufacturers are not likely to produce and implement

solutions for these “poor markets”. On the other hand, DiffServ is supported by barely any network device nowadays, and DiffServ routers can be developed with few resources. Hence, DiffServ is a feasible solution.

- It is not efficient that traffic occupies several links in a chain for eventual discarding. Similarly, it is not efficient that packets enter a node, enter the egress queue, generate delay to other packets, and are eventually discarded.
- The limitations on delay imposed by the voice traffic require that all nodes conduct traffic shaping in order to prevent links from operating under saturation conditions.

Hence,

- Edge nodes must apply ingress policing, which means that traffic exceeding certain thresholds must be discarded before it can cause damage to legitimate traffic.
- All nodes must know the maximum capacity of the links under unsaturated conditions, and egress queuing disciplines must perform traffic shaping.
- The available bandwidth must be assigned differently for the three aggregated traffic classes:
 - Signaling must have maximum priority, with a long queue to prevent packet-drop and a share of 1% of the total bandwidth.
 - Voice must have priority on data and shall be assigned a guarantee corresponding to the maximum occupation of telephony channels.
 - Data must have the rest of all available resources, with minimum priority.

The lecturer may use the TUCAN3G deliverable TUCAN3G-D51 (2014), especially Chapter 6, for exhaustive explanations regarding the most appropriate solution to the case study.

4. HOMEWORK ACTIVITY

Students should be organized in groups of 3-4 student. It is advisable, though not necessary, that two of them have a laptop computer available with Linux or alternately with a Live-linux pendrive. See Ubuntu (2014) for guidance on how to use it.

STEP 1.- Understand how advanced queuing disciplines behave. Students must visit the LARTC (2014) site and read about queuing disciplines in general, and particularly on HTB, TBF and SFQ qdiscs. The students must also understand how filters work, especially DSCP filtering. If the students have access to a Linux system, it is a good idea to test the different qdiscs, individually and combined, with the assistance of a traffic injector like (D-ITG, 2014). However, this experimental part is not obligatory. Either after theoretical study or after testing, the students must describe the behaviour of each queuing discipline.

STEP 2.- Formally define the types of traffic that the network must differentiate between, how priorities should be assigned, and how and where traffic is differentiated. What DSCP values would you assign to each traffic class? Why? Where should packets be marked? BONUS: how should filters be configured in queuing disciplines in order to recognise the different DSCP marks?

STEP 3.- Take the transport network architecture illustrated in Figure 4, and apply it to a segment of the Napo network. Consider that each village between Santa Clotilde and Tuta Pishco has a 3G small cell, except Santa Clotilde, which has two cells and consider also that the gateway is connected to a satellite communications system in Negro Urco. Design the specific network architecture for the network described. Identify all systems and links, and describe each system's functions.

STEP 4.- Consider that each HNB generates a maximum voice traffic throughput of 320 kbps in the uplink and 960 kbps in the downlink. It also requires a capacity of 500 kbps in the uplink and 2500 kbps in the downlink for data traffic (although the data traffic may be higher if the available capacity allows it), and the signalling traffic is always below 1% of the rest of the traffic in each direction. The delay for signalling and for voice must be under 50 ms, and the packet loss must be 0% for signalling and <1% for voice traffic. Links generate a delay of 5 ms each. Calculate what capacity is required for each link in order to bear all the traffic properly.

STEP 5.- Policing: the transport network is going to accept a limited amount of traffic from/to each HNB. Traffic exceeding these maximums must be discarded in order to prevent the transport network becoming saturated. Say exactly where these limitations need to be imposed and what the limits imposed in each point are.

STEP 6.- Traffic shaping and bandwidth sharing: Use either Hierarchical Token Bucket (HTB) or PRIO queuing disciplines for the root queue in the egress block of each router. Both queuing disciplines have classes, so they have several bands and choose what child queuing discipline manages the traffic in each band. Consider only three possibilities for the child qdiscs: FIFO, TBF and SFQ. The traffic offered to each link must be limited to the maximum previously calculated. How much is this limitation for each system and how can this be done in the proposed queuing disciplines? What is the share that each traffic class has in each system? What are the differences between HTB and PRIO in the way they support both the traffic shaping and the bandwidth sharing? What child queuing disciplines should be used in each case?

STEP 7.- Put everything together and propose a complete configuration for a correct QoS management in the network. In the proposal, should IP packets be marked before entering the transport network? Can packets keep their marks after abandoning the transport network?

With the answer to all the previous questions, compose a technical proposal on how the transport network may be implemented.

Solution and evaluation criteria

The network has four routers in a row, with wireless links connecting each hop. Each wireless link needs to have a capacity higher than the maximum throughput expected across that link. The link starting in Santa Clotilde only bears the traffic exchanged by the two HNB in that town, but the next one must add the traffic exchanged with Tachsa's traffic. The link from Negro Urco to Tuta Pishco only bears Tuta's traffic. Finally, the gateway in Negro Urco aggregates all the traffic at the interface to the satellite link.

The policing must be done in the ingress blocks to which HNB are connected. Packets discarded later will impact negatively on the network's performance. Then, each egress block must use a traffic shaper to limit the traffic to the maximum expected in the design. HTB allows this maximum in the root qdisc to be imposed, while PRIO does not. Hence, in the case of using PRIO the traffic in child queues must be shaped using TBF. In HTB, the bands for voice and signalling must be dimensioned for 100% of traffic expected, and the rest must be assigned to data. Child queues for voice and signalling should be FIFO, but SFQ should be used for data traffic.

Regarding the traffic marking and unmarking, HNB may mark the DSCP of each packet (this is what actually happens in the reality). Alternatively the edge node may classify the traffic, but this is a difficult task. DSCP are referred to as hexadecimal or decimal numbers that represent the 6 highest bits in the DS field. Hence, filters must read the DS byte, do a two-bit shift (introducing zeros on the two left-most bits) and interpret the value obtained.

Last but not least, the students must understand the importance of unmarking the DSCP in all packets in the gateway for the outgoing traffic. Packets with non-zero DSCP may be interpreted differently out of the transport network. The semantics of the DSCP are internal to the DiffServ domain.

The lecturer may use the TUCAN3G deliverable TUCAN3G-D52 (2014), particularly 8.2, for full explanations regarding the most appropriate solution to the case study problem.

BIBLIOGRAPHY

D-ITG, Available from: <<http://traffic.comics.unina.it/software/ITG/>> (Last visit: July 2014).

ECONOMIC COMMISSION FOR LATIN AMERICA AND THE CARIBBEAN (2012) Anuario Estadístico de América Latina y el Caribe 2012. Santiago de Chile. ECLAC.

INTERNET ENGINEERING TASK FORCE (1998) RFC 2475: An Architecture for Differentiated Services. IETF.

INTERNET ENGINEERING TASK FORCE (1998) RFC 2474: Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPV6 Headers. IETF.

INTERNET ENGINEERING TASK FORCE (1999) RFC 2597: Assured Forwarding PHB Group. IETF.

INTERNATIONAL TELECOMMUNICATIONS UNION (2001) Recommendation ITU-T G-1010: End-user multimedia QoS categories. Geneva. ITU-T.

INTERNET ENGINEERING TASK FORCE (2002) RFC 3246: An Expedited Forwarding PHB (Per-Hop Behavior). IETF.

INTERNATIONAL TELECOMMUNICATIONS UNION (2008) Recommendation ITU-T E-800: Quality of telecommunication services: concepts, models, objectives and dependability planning – Terms and definitions related to the quality of telecommunication services. Geneva. ITU-T.

Kilikki K 1999, Differentiated services for the Internet. Macmillan Technical Publishing.

Klein G et al. 2007, Diagnóstico Regional en Materia de Políticas, Estrategias y Regulaciones de las Telecomunicaciones y Desarrollo de las Tecnologías de la Comunicación e Información (TIC). Aplicaciones en las áreas de Educación, Salud y Catástrofes en Zonas Rurales, Regiones Aisladas y Urbano-Marginales: *Final Report ITU*, Brasilia.

LARTC, Available from: <<http://www.lartc.org>> (Last visit: July 2014).

TUCAN3G (2013) Wireless technologies for isolated rural communities in developing countries based on cellular 3G femtocell deployments. Project funded by FP7, European Commission. Available from: <<http://www.ict-tucan3g.eu/>>.

TUCAN3G-D21 (2013) Socio-economic scenarios, technical specifications and architecture for the proof of concept. Deliverable D21 of TUCAN3G project. Available from: <http://www.ict-tucan3g.eu/index.php?option=com_ckforms&view=ckforms&id=4&Itemid=88&doct=deliverables&doc=TUCAN3G_D21EHASd.pdf> (June 2013).

TUCAN3G-D51 (2013) Technical requirements and evaluation of WiLD, WIMAX and VSAT for backhauling rural femtocells networks. Deliverable D51 of TUCAN3G project. Available from: <http://www.ict-tucan3g.eu/index.php?option=com_ckforms&view=ckforms&id=4&Itemid=88&doct=deliverables&doc=TUCAN3G_D51URJCa.pdf> (October 2013).

TUCAN3G-D52 (2014) Heterogeneous transport network testbed deployed and validated in laboratory. Deliverable D52 of TUCAN3G project. Available from: <http://www.ict-tucan3g.eu/index.php?option=com_spcom&Itemid=88> (October 2014).

Ubuntu. Available from: <<http://www.ubuntu.com>> (Last visit: July 2014).

UNITED NATIONS DEVELOPMENT PROGRAMME (2013) Human Development Report 2013: The raise of the South: Human progress in a diverse world. New York. United Nations Development Programme.



GDEE

GLOBAL
DIMENSION IN
ENGINEERING
EDUCATION

<http://www.gdee.eu>



This project is funded by

