



TD(05)036

4th Management Committee and Technical Meeting

Wuerzburg, October 13-14th, 2005

Implementation of VoIP over Satellite

Ulla Birnbacher, Ronald Toegl
Institute of Communication Networks and Satellite Communications, Graz University of Technology
Inffeldgasse 12, 8010 Graz
Ulla.birnbacher@tugraz.at

Abstract

IP Telephony is one of the fastest growing technologies. It is based on the Internet protocol suite. However to provide the same services and service qualities like legacy telephony networks, special care has to be taken for the IP network design.

The Internet protocol suite is also increasingly used for satellite communications, which offers wide-area broadband connections to any point within the satellite's footprint. In order to implement Voice over IP (VoIP) over satellite links, the network design has additionally to consider the physical characteristics of the link.

This paper presents the implementation of VoIP over satellite links, which are also used in parallel for data transfer and video conferencing. Quality of Service (QoS) is thereby the key technology to enable the combination of data and voice services. It is investigated which level of service quality can be achieved both on the LAN and on the Satellite Link by using QoS mechanisms available in off-the-shelf routers and switches. This paper includes a presentation of the basic technologies, the architectural design and an in-depth discussion of the configuration used. Finally, the resulting service quality is discussed.

Keywords

VoIP, Satellite Communication, QoS

Working Group 1

IMPLEMENTATION OF VOIP OVER SATELLITE

Ulla Birnbacher, Ronald Toegl
Institute of Communication Networks and Satellite Communications,
Graz University of Technology, Graz, Austria

Abstract— IP Telephony is one of the fastest growing technologies. It is based on the Internet protocol suite. However to provide the same services and service qualities like legacy telephony networks, special care has to be taken for the IP network design.

The Internet protocol suite is also increasingly used for satellite communications, which offers wide-area broadband connections to any point within the satellite's footprint. In order to implement Voice over IP (VoIP) over satellite links, the network design has additionally to consider the physical characteristics of the link.

This paper presents the implementation of VoIP over satellite links, which are also used in parallel for data transfer and video conferencing. Quality of Service (QoS) is thereby the key technology to enable the combination of data and voice services. It is investigated which level of service quality can be achieved both on the LAN and on the Satellite Link by using QoS mechanisms available in off-the-shelf routers and switches. This paper includes a presentation of the basic technologies, the architectural design and an in-depth discussion of the configuration used. Finally, the resulting service quality is discussed.

Index Terms—DiffServ, Satellite Communication, QoS, VoIP.

I. INTRODUCTION

Data networks based on the IP protocol are capable to transport nearly any type of service, starting from simple data transfer to real-time services like video conferencing and telephony services. Currently, IP Telephony is one of the fastest growing technologies. However to provide the same services and service qualities like legacy telephony networks, special care has to be taken for the IP network design. Local area networks (LANs) offering hundreds or thousands of MBit/s provide enough bandwidth to support delay and jitter sensitive services like Voice over IP (VoIP). The required Quality of Service (QoS) is thereby enabled by simple bandwidth over-provisioning and keeping the network load low.

In wireless networks and satellite communication networks in particular, bandwidth is a scarce resource and has to be

used very efficiently. Internet protocols are also increasingly used for satellite communications, which offer wide-area broadband connections to any point within the satellites' footprint. In order to implement VoIP over satellite links, network design has to provide the required QoS by other means than bandwidth over-provisioning, taking also into account the physical characteristics of the link. Off-the-shelf network components like routers and switches are already supporting QoS technologies (e.g. DiffServ, evaluation of ToS field, priority queuing).

At the Institute of Communication Networks and Satellite Communications of Graz University of Technology several satellite ground stations with indoor and outdoor equipment and Cisco networking devices are available for testing. Currently research focuses on the implementation of all-IP networks over satellite. QoS is thereby the key technology to enable the combination of data and voice network services.

This paper presents the implementation of VoIP by using QoS mechanisms readily available in standard Cisco routers and switches. The satellite link is also used in parallel for data transfer and video conferencing. It is investigated, if the perceived quality of voice and video communication can be increased by employing IP QoS mechanisms. After presenting the basic technologies, the architectural design and an in-depth discussion of the configuration used, the resulting service quality is discussed finally.

II. LAYER 2 CHARACTERISTICS AND SERVICES REQUIREMENTS

To carry voice over different network technologies, their characteristics need to be considered. Legacy Public Switched Telephony Networks (PSTN) are centrally controlled networks which establish circuit switched connections between simple and cheap end-user devices. In modern telephony networks, a packet-switched backbone is used to carry the voice digitally. These networks are designed to provide a guaranteed voice service quality.

Local Area Networks (LAN), like Ethernet follow a completely different approach. They offer plenty of bandwidth

(10, 100, 1000 or more Mbit/s) to all hosts attached but no central control. Modern Ethernet networks, based on twisted pair or fibre optics add full duplex connections without the risk of collisions. By using switches, packet-switched point-to-point connections are available. On switched networks, it is possible to provide certain qualities of service, if care is taken in the network design.

By the use of communication satellites, networks can be interconnected on a global scale, also in geographic regions with poor terrestrial backbone infrastructure. Satellites in the geostationary orbit (36.000 km above the Earth's surface) are characterized by the high propagation delay, but can provide large geographic areas with communication services.

A specific use of communication satellites is the interconnection of LANs. Two or more geographically separated LANs can be interconnected over satellite links using the IP protocol. By implementing IP on satellite links, a variety of applications can be supported. Since the network protocol is the same as on LANs and the Internet, the full diversity of data services, video conferences and VoIP can be applied on satellite links as well. This flexibility allows the use of a single satellite link for any kind of application. In [5] and [6] it is shown that it is feasible to carry VoIP data over satellite links.

The next step is to allow the simultaneous operation of various services. To prevent mutual interferences, the application of QoS technologies is of key importance.

III. NETWORK ARCHITECTURE

The network consists of two separate LANs, connected by a satellite link. Base station 1 (BST1) is the main site. It connects to the Internet and to the ISDN network. It is therefore home to the VoIP gateway and also to a gatekeeper, a Cisco CallManager Express (CME) 3.0. The remote site (base station 2 - BST2) has similar equipment, but no additional network links. Both feature the following basic components: A Cisco router with Voice support (Cisco 1760V and Cisco 2801), a managed switch, one Cisco IP Phone, a notebook computer and a Polyspan video conference unit. This setup provides a variety of services, including but not restricted to IP telephony, video conferences, Web, FTP and e-mail clients, FTP and NTP servers and more.

An overview is provided in figure 1. The two sites are connected by two 1 Mbit/s satellite links realising a FDM network, using the geostationary SESAT satellite positioned at 36 degrees East. The satellite ground stations consist of a parabolic satellite antenna with 1.5m and 2.4m diameter respectively, connected to an L-Band up-/down-converter and to a Radyne Comstream satellite modem (QPSK modulation, Viterbi Forward Error Correction with code rate $\frac{3}{4}$). The

Cisco routers are directly connected to the RS.449 interface of the satellite modem.

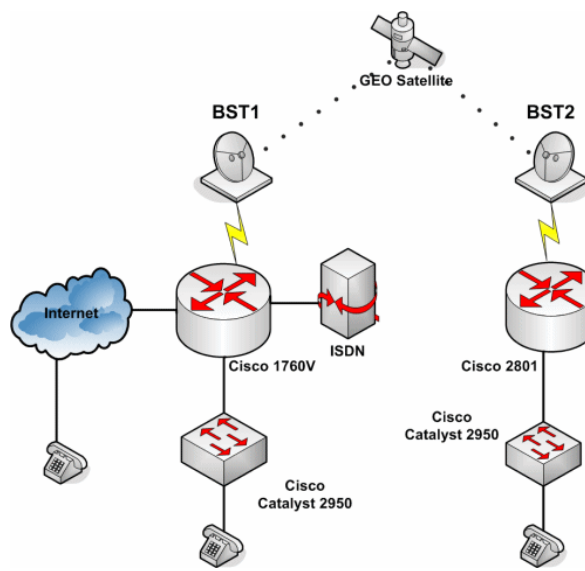


Figure 1: Architectural overview

The VoIP implementation is based on Cisco's solution for IP telephony. In our setup, Cisco Call-Manager Express (CME) software which provides all common telephony services is used. It is based on Cisco IOS - it therefore can be operated directly on the routers, reducing the hardware and maintenance needs. It is designed to serve a few dozens of users, especially in small single deployments or at small remote sites. At both base stations a CME acts as gatekeeper. The routers on which the CME software is operating, provide gateway services at the same time. Both sites can operate a local IP telephony network on their own, but inter-CallManager connections also allow conducting calls to the other site. In addition, BST1 provides a gateway to the PSTN via the PBX of Graz University of Technology. Since both sites feature a full CME, the setup can be characterised as multi-site WAN with distributed call processing.

The IP phones communicate with their local CallManager with the Cisco proprietary Skinny Client Control Protocol. The voice data is carried by the Real-time Transport Protocol (RTP): For local calls, where plenty of bandwidth is available, the G.711 codec is used, while remote calls use G.729. Another difference between local and remote calls is the use of H.323 protocols for inter-CallManager communication, for instance the use of the H.225 protocol for establishing remote calls.

Whereas this setup works fine, when there is enough bandwidth (over)provisioned, it will fail to provide acceptable voice quality when the satellite link is saturated. The key is the application of QoS mechanisms.

IV. QoS

In the context of computer networks, Quality of Service (QoS) of traffic is defined as providing guarantees on certain boundaries on four properties of the traffic

- Reliability
- Bandwidth
- Delay
- Jitter.

These definitions allow us to name the QoS requirements and common recommendations for voice traffic [2].

- Voice traffic should be marked to the DSCP value of Expedited Forwarding (EF).
- Loss should be less than 1 percent. At this rate, it will not be noticeable to the user.
- According to the ITU G.114 standard, one-way, end-to-end, i.e. mouth-to-ear, delay should not exceed 150 ms. It is however recognised that this is not always achievable due to the delays on satellite links.
- Average one-way jitter should be targeted at less than 30 ms.
- Enough guaranteed priority bandwidth should be available for each call.

These goals are achieved both in the LAN and on the satellite link by using Differentiated Services. Traffic is not separated in flows but rather in classes. Each packet carries a description of its class and is therefore treated differently in every router along the way. DiffServ is configured locally at every hop in the network. All packets in a class share the same resources, but exclusively so. Different classes will be assigned different resources and therefore will not interfere with each other.

Differentiated Services can be applied in switches and routers. Each frame or packet needs to carry its classification information. Layer 2 802.1q frames may carry a Class of Service (CoS) field and the layer 3 IP header provides the ToS byte, which contains among others a 3-bit IP Precedence field. Besides pure bits, the interpretation of them is more important.

Traffic Type	Layer 2 CoS	IP Precedence	DSCP	PHB
Voice RTP	5	5	46	EF
Voice control	3	3	24	CS3
Video conferencing	4	4	34	AF41
Data	0, 1, 2	0, 1, 2	10 to 22	BE to AF23

Table 1 IP Telephony traffic Mapping guidelines [4]

In the DiffServ Architecture (RFC 2475) the ToS byte of the IP header was redefined (RFC 2474). It now features a 6-bit Differentiated Services Code Point (DSCP) descriptor. DSCP is the basis for the various DiffServ actions. Certain values are assigned to special per-hop behaviours. The actual QoS requirements for each single traffic class's user data need to be taken into account. Table 1 shows in [4] suggested mappings for voice and video applications. Defining a mapping of CoS and IP Precedence values to DSCP is necessary, because several different interpretations exist and many software applications do not support DSCP yet.

However, setting these DSCPs alone does not result in any per-hop behaviours unless the network nodes are specifically configured to do so. These behaviours are dominated by two mechanisms. *Queuing* is the process of assigning the different classes of traffic to their queues. This separation of classes makes sure that they do not interfere with each other too much. Once the queues are being filled, they need to be served. How this is done is controlled by *Scheduling*. The simplest scheduling scheme is the sequential round-robin fair-queuing. When the transmitting queue becomes idle, the next one is served and one packet transmitted. An extension, weighted round robin (WRR), performs better if the packets in queues have different lengths. Here a ratio of bandwidth is assigned to the non-empty queues and scheduling done accordingly. A simple example: if given a queue A with weight 1 and queue B weight 3, then 75% of the bandwidth is assigned to the traffic in B. The so called Class-Based Weighted Fair Queuing (CBWFQ) scheme integrates these mechanisms. Traffic classes are assigned to queues, and the queues are being served according to the WRR settings.

The CBWFQ scheme provides separation of different classes and allows prioritisation. However to provide voice a fragmentation-free, low loss, low latency link to achieve toll quality, a more powerful scheme is needed. The solution found in today's network devices is the priority queue. It has strict priority: no other queues are served, until they are empty. To prevent other traffic from starvation, the priority queue needs to be policed not to exceed a given maximum bandwidth. A CBWFQ scheme combined with a priority queue is referred to as Low Latency Queuing (LLQ).

In our setup each Cisco telephony device (CallManager, IP phone) sets correct CoS and DSCP markings for its frames and packets conform to table 1. A multi-layered approach is taken to enable QoS guarantees throughout the network. At the access layer, the Catalyst 2950 switches perform a first classification. Of course the switches also apply per-hop behaviours. The LLQ implementation and configuration is illustrated in figure 2. Absolute priority is given to voice traffic which carries the DSCP for Expedited Forwarding (RFC 3246) and is assigned the priority queue. Real-time video and call control data are treated with CBWFQ in a

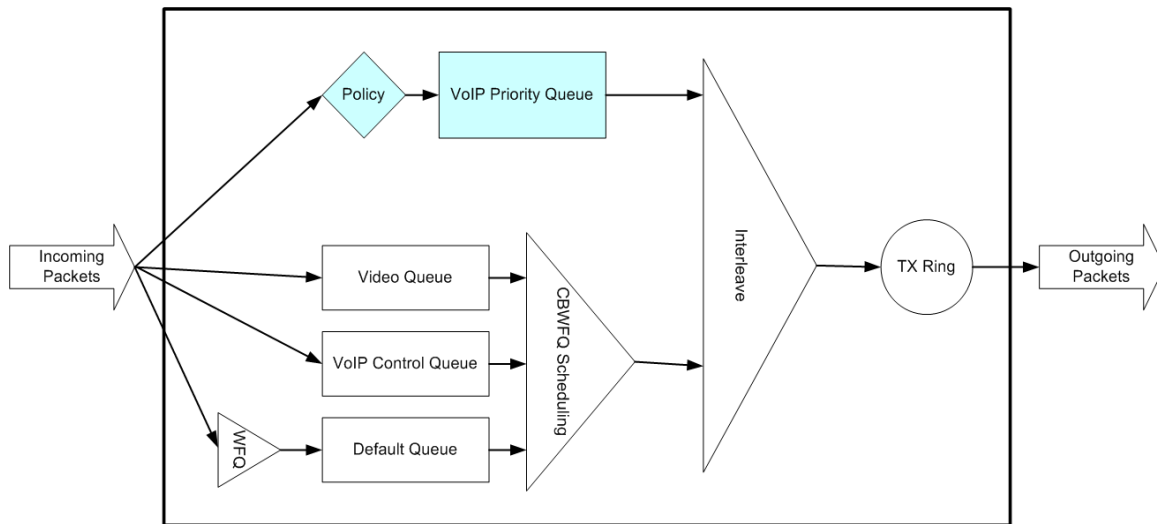


Figure 2 LLQ scheme

simplified, Assured Forwarding (RFC 2597)-like scheme. On the satellite link, where resources are especially scarce, bandwidth guarantees are given to real-time and control data. The highest priority class, which is destined for voice is assigned a guaranteed bandwidth of 160 kbit/s. This is enough for at least four G.729 calls at the same time.

Voice control is not as sensitive, and it is assigned a guaranteed bandwidth of 8 kbit/s. For video conference traffic nearly the maximum bandwidth is reserved. It is not guaranteed the same low latency as for voice, but the traffic may exceed the guaranteed limit without being dropped. When there is no need for the complete reserved amount of bandwidth, other traffic is allowed. The sum of all reservations may not exceed the maximum available link speed minus a reserve of 10% for network control traffic. Finally, the remaining best-effort traffic is allowed to take all of the free bandwidth. This is done using weighted fair-queuing (WFQ), so all streams are allowed to transmit from time to time.

Once the packets are received on the other site, layer 2 information is generated from the DSCP field prior to sending the packet out to the switched network.

V. RESULTS

Voice quality is determined by QoS conditions. To evaluate the QoS mechanism, the following setup was active. Only the CallManager of BST1 was used, to allow monitoring the RTP voice stream on the WAN link.

For a first functionality test under lab conditions, the satellite link was replaced by a wired link between the satellite modems at intermediate frequency. The bandwidth available

was 2 Mbit/s. The complete traffic was monitored with Ethereal [3], a free, but powerful network sniffer. To challenge the QoS mechanism, four types of traffic were present: A bidirectional video conference using the adaptive H.263 codec at up to 768 kbit/s. One voice call from the remote to the local site, using the G.711 uncompressed codec at 64 kbit/s. Voice control traffic is based on Cisco's Skinny protocol. And finally, several FTP streams in both directions to saturate the link and simulate full network load. With QoS disabled, neither video or voice performance was acceptable.

Once QoS was activated, excellent voice and video performance was experienced.

For a more detailed analysis, the complete traffic was monitored at the trunk line between the switch and the router at BST1. A graph of the different traffic classes was created with Ethereal (figure 3). The plot shows a sample run for about 2 minutes under full load. It is clearly visible how low priority traffic (c – data) yields to high (b - video) and highest (d – voice) priorities. Line a shows the overall duplex traffic on the link - it is always at the maximum value of 2 Mbit/s. In the beginning, there is FTP and video traffic. The video codec is adaptive, and does not need much bandwidth when there is no movement in front of the camera. In the first 5 seconds, FTP traffic is allowed to use almost all of the bandwidth. Then increased movement activity was staged before the camera, causing a rise in the bandwidth requirement of the video transmission. Now FTP traffic is not permitted as much bandwidth any more and slowed down. When the video is not transmitting any movements (at about 20s) it picks up speed again. At about 60 seconds a very short VoIP call is performed (line d). At once the FTP traffic is slowed down to allow voice communication. Observe that the G.711 codec is non-adaptive and always requires a constant bandwidth. Ten

seconds later, a second voice call is started. At the same time, some movements were transmitted by the video stream. Again, it is the FTP traffic that yields to traffic with higher priority. The voice traffic in this experiment has the following characteristics. It requires a bandwidth of 80 kbit/s (64 kbit/s for the codec and 16 kbit/s for Ethernet/IP/UDP/RTP

reservation on the satellite link. The setup could be extended to other media (802.11 WLAN, 802.16 WiMax, Free Space Optics) while maintaining QoS guarantees. Finally, security should be reviewed and improved. A firewall and VPNs could be key components. A specific focus could be the prevention of eavesdropping the broadcast satellite link.




Figure 3: Traffic analysis. a - Overall traffic. b - Adaptive video traffic. c - FTP file transfer. d - Voice

headers). The maximum delay per packet is 42.55 [ms] and the maximum jitter is 12.67 [ms]. No packets were lost.

In a second test series, the sites were connected via satellite links. The round-trip-time was 548 [ms], measured by a “ping” between both routers. A local-notebook-to-remote-notebook ping resulted in a RTT of 615 ms. The used Cisco IP phones feature an internal collection of statistical data regarding the voice connections. The collected statistics retrieved a maximum encountered jitter of 68 ms. The average jitter on the RTP connection between gateway and phone was only 4.75 ms though. The highest packet loss rate (including packets dropped within the telephone, not only those dropped on the link) was a mere 2 out of 2570 packets. It can be concluded, that the QoS requirements for VoIP were met. The empirical voice quality evaluation suggested very good results, no matter if the link was saturated with other traffic or not.

VI. CONCLUSION AND FUTURE PLANS

With the described setup, a flexible and mobile communication platform using satellite links has been created. By using the Internet protocol stack, any standard application running over the (terrestrial) Internet can be used. The implementation of VoIP shows that Quality of Service (QoS) is required for enabling the convergence of data and voice networks. By using the available QoS mechanisms of off-the-shelf routers and switches excellent service quality can be achieved and furthermore allows for good utilisation of scarce satellite resources and provision of several simultaneous services.

Future extensions could include the introduction of advanced QoS using call admission control and resource

REFERENCES

- [1] Jonathan Davidson, James Peters, and Brian Grace. *Voice over IP Fundamentals*. Indianapolis: Cisco Press, 2000.
- [2] Tim Szigeti and Christina Hattingh. *End-to-End QoS Network Design*. Indianapolis: Cisco Press, 2005.
- [3] Ethereal <http://www.ethereal.com/>
- [4] *Cisco IP Telephony*, Solution Reference Network Design, 2004, <http://www.cisco.com/>
- [5] *Voice over IP over satellite links*, Cruickshank, H., Sanchez, A., Zhili Sun and Carro, B.; Electronics, Circuits and Systems, 2001. ICECS 2001. The 8th IEEE International Conference on Volume 1, 2-5 Sept. 2001 Page(s):473 - 476 vol.1
- [6] *Voice over IP service and performance in satellite networks*, Nguyen, T., Yegenoglu, F., Sciuto, A., and Subbarayan, R., Communications Magazine, IEEE Volume 39, Issue 3, March 2001 Page(s):164 - 171