

Voice over Internet Protocol (VOIP) - Future of communication

Mrs. Roopali garg*, Jagpreet Sidhu**

*roopali.garg@gmail.com, **jagpreetsidhu@gmail.com

Abstract - During the recent Internet stock bubble, articles in the trade press frequently said that, in the near future, telephone traffic would be just another application running over the Internet. Such statements gloss over many engineering details that preclude voice from being just another Internet application. This paper deals with the technical aspects of implementing voice over Internet protocol (VoIP), without speculating on the timetable for convergence. First, the paper discusses the factors involved in making a high-quality VoIP call and the engineering tradeoffs that must be made between delay and the efficient use of bandwidth. After a discussion of codec selection and the delay budget, there is a discussion of various techniques to achieve network quality of service. Since call setup is very important, the paper next gives an overview of telephony routing over IP (TRIP). Finally, the paper explains some VoIP issues with network address translation and firewalls.

Keywords—H.323, Internet telephony, MGCP, SIP, telephony routing over IP (TRIP), voice over IP (VoIP), voice quality.

I. INTRODUCTION




Fig. 1. Business use of VoIP.

There is a plethora of published papers describing various ways in which voice and data communications networks may “converge” into a single global communications network. This paper deals with the technical aspects of implementing VoIP, without speculating on the timetable for convergence. A large number of factors are involved in making a high-quality VoIP call. These factors include the speech codec, packetization, packet loss, delay, delay variation, and the network architecture to provide QoS. Other factors involved in making a successful VoIP call include the call setup signaling protocol, call admission control, security concerns, and the ability to traverse NAT and firewall. Although VoIP involves the transmission of digitized voice in packets, the telephone itself may be analog or digital. The voice may be digitized and encoded either before or concurrently with packetization. Fig. 1 shows a business in which a PBX is connected to VoIP gateway as well as to the

local tele- phone company central office. The VoIP gateway allows tele- phone calls to be completed through the IP network. Local calls can still be completed through the telephone company as in the past. The business may use the IP network to make all calls between its VoIP gateway connected sites or it may choose to split the traffic between the IP network and the PSTN based on a least-cost routing algorithms configured in the PBX. VoIP calls are not restricted to telephones served directly by the IP network. We refer to VoIP calls to telephones served by the PSTN as “off-net” calls. Off-net calls may be routed over the IP network to a VoIP/PSTN gateway near the destination telephone.

An alternative VoIP implementation uses IP phones and does not rely on a standard PBX. Fig. 2 is a simplified diagram of an IP telephone system connected to a wide area IP network. IP phones are connected to a LAN. Voice calls can be made locally over the LAN. The IP phones include




Fig. 2. VoIP from end to end.

Table 1 Characteristics of Several Voice Codecs

Codec	Algorithm	Frame Size/ Lookahead	Usual Rate	Comments
G.711	PCM	0.125 ms/0	64 Kb/s	Universal use
G.722		0.125 ms/1.5 ms	48, 56 or 64 Kb/s	Wideband coder
G.726	ADPCM	0.125 ms/0	32 Kb/s	High quality, low complexity
G.728	LD-CELP	0.625 ms/0	16 Kb/s	High quality in tandem; Recommended for cable
G.729(A)	CS-ACELP	10 ms/5 ms	8 Kb/s	Widespread use
G.729	Hybrid CELP	10 ms/5 ms	11.8 Kb/s	High quality/complexity; Recommended for cable
G.723.1(6.3)	MPC-MLQ	30 ms/7.5 ms	6.3 Kb/s	Video conferencing origin
G.723.1(5.3)	ACELP	30 ms/7.5 ms	5.3 Kb/s	Video conferencing origin
IS-127	RCELP	20 ms/5ms	Var. 4.2 Kb/s avg.	
AMR	ACELP	20 ms	Var. 4.75-12.2 Kb	Compatible w. No. Amer. & Japanese digital cellular, WCDMA (not CDMA2000); Nokia IPR

codecs that digitize and encode (as well as decode) the speech. The IP phones also packetize and depacketize the encoded speech. Calls between different sites can be made over the wide area IP network. Proxy servers perform IP phone registration and coordinate call signaling, especially between sites. Connections to the PSTN can be made through VoIP gateways.

II. VOICE QUALITY

Many factors determine voice quality, including the choice of codec, echo control, packet loss, delay, delay variation (jitter), and the design of the network. Packet loss causes voice clipping and skips. Some codec algorithms can correct for some lost voice packets. Typically, only a single packet can be lost during a short period for the codec correction algorithms to be effective. If the end-to-end delay becomes too long, the conversation begins to sound like two parties talking on a Citizens Band radio. A buffer in the receiving device always compensates for jitter (delay variation). If the delay variation exceeds the size of the jitter buffer, there will be buffer overruns at the receiving end, with the same effect as packet loss anywhere else in the transmission path. For many years, the PSTN operated strictly with the ITU standard G.711. However, in a packet communications network, as well as in wireless mobile networks, other codecs will also be used. Telephones or gateways involved in setting up a call will be able to negotiate which codec to use from among a small working set of codecs that they support. Codecs: There are many codecs available for digitizing speech. Table 1 gives some of the characteristics of a few standard codecs.1




Fig. 3. Effect of codec concatenation on an MOS.

The quality of a voice call through a codec is often measured by subjective testing under controlled conditions using a large number of listeners to determine an MOS. Several characteristics can be measured by varying the test conditions. Important characteristics include the effect of environmental noise, the effect of channel degradation (such as packet loss), and the effect of tandem encoding/decoding when interworking with other wireless and terrestrial transport networks. The latter characteristic is especially important since VoIP networks will have to interwork with switched circuit networks and wireless networks using different codecs for many years. The general order of the fixed-rate codecs listed in the table, from best to worst performance in tandem, is G.711, G.726, G.729e, G.728, G.729, G.723.1. Quantitative results are given in [1]. Since voice

quality suffers when placing low-bit-rate codecs in tandem in the transmission path, the network design should strive to avoid tandem codecs whenever and wherever possible.

Concatenation and Transcoding: The best packet network design codes the speech once near the speaker and decodes it once near the listener. Concatenation of low-bit-rate speech codecs, as well as the transcoding of speech in the middle of the transmission path, degrades speech quality. Fig. 3 shows the MOSs of several codecs with and without concatenation. (These results are from [1]. An MOS of 5 is excellent, 4 is good, 3 is fair, 2 is poor, and 1 is very bad. Note that G.729 \times 2 means that speech coded with G.729 was decoded and then recoded with G.729 before reaching the final decoder. G.729 \times 3 means that three G.729 codecs were concatenated in the speech path between the speaker and listener.) Fig. 4 shows the MOSs resulting from the interworking of different codecs, possibly in a transcoding situation.

III. TRANSPORT

Typical Internet applications use TCP/IP, whereas VoIP uses RTP/UDP/IP. Although IP is a connectionless best effort network communications protocol, TCP is a reliable transport protocol that uses acknowledgments and retransmission to ensure packet receipt. Used together, TCP/IP is a reliable connection-oriented network communications protocol suite. TCP has a rate adjustment feature that increases the transmission rate when the network is uncongested, but quickly reduces the transmission rate when the originating host does not receive positive acknowledgments from the destination host. TCP/IP is not suitable for real-time communications, such as speech transmission, because the acknowledgment/retransmission feature would lead to excessive delays. UDP provides unreliable connectionless delivery service using IP to transport messages between end points in an internet. RTP, used in conjunction with UDP, provides end-to-end network transport functions for applications transmitting real-time data, such as audio and video, over unicast and multicast network services.[2] RTP does not reserve resources and does not guarantee quality of service. A companion protocol RTCP does allow monitoring of a link, but most VoIP applications offer a continuous stream of RTP/UDP/IP packets without regard to packet loss or delay in reaching the receiver.

Although transmission may be inexpensive on major routes, in some parts of the world as well as in many private networks, transmission facilities are expensive enough to merit an effort to use bandwidth efficiently. This effort starts with the use of speech compression codecs. Use of low bandwidth leads to a long packetization delay and the most complex codecs. An engineering tradeoff must be made to achieve an acceptable packetization delay, an

acceptable level of codec complexity, and an acceptable call transmission capacity requirement.

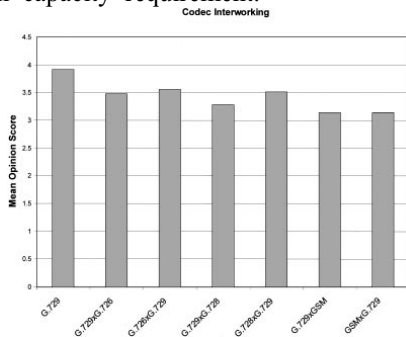


Fig. 4. Effects of transcoding.

Another technique for increasing bandwidth efficiency is voice activity detection and silence suppression. Voice quality can be maintained while using silence suppression if the receiving codec inserts a carefully designed comfort noise during each silence period. For example, Annex B of ITU-T Recommendation G.729 defines a robust voice activity detector that measures the changes over time of the background noise and sends, at a low rate, enough information to the receiver to generate comfort noise that has the perceptual characteristics of the background noise at the sending telephone [3].

Coding and packetization result in delays greater than users typically experience in terrestrial switched circuit networks. As we have seen, standard speech codecs are available for output coding rates in the approximate range of 64 to 5 kb/s. Generally, the lower the output rate, the more complex the codec. Packet design involves a tradeoff between payload efficiency (payload/total packet size) and packetization delay (the time required to fill the packet). For IPv4, the RTP/UDP/IP header is 40 bytes. A payload of 40 bytes would mean 50% payload efficiency. At 64 kb/s, it only takes 5 ms to accumulate 40 bytes, but at 8 kb/s it takes 40 ms to accumulate 40 bytes. A packetization delay of 40 ms is significant, and many VoIP systems use 20-ms packets despite the low payload efficiency when using low-bit-rate codecs. For continuous speech, the call transmission capacity requirement BW (in kb/s) is related to the header size H (in bits), the codec output rate R (in kb/s) and the payload sample size S (in milliseconds) as

$$BW = R + H/S$$

Fig. 5 shows a plot of BW versus R and S assuming H-320b. There are several header compression algorithms that will improve payload efficiency [4]–[6]. The 40-byte RTP/UDP/IP header can be compressed to 2–7 bytes. A typical compressed header is four bytes, including a two-byte checksum. In an IP network, header compression must be done on a link-by-link basis, because the header must be restored before a router can choose an outgoing interface. Therefore, this technique is most suitable for

low-speed access links. Fig. 6 shows a plot of BW versus R and S assuming H-320b

The lowest BW requirements lead to a long packetization delay and the most complex codecs. An engineering tradeoff must be made to achieve an acceptable packetization delay, an acceptable codec complexity, and an acceptable call bandwidth requirement. The following sections discuss quality and bandwidth efficiency in more detail.




Fig. 5. The varying bands, from top to bottom, represent the following VoIP bandwidth requirements (40-byte headers): 120–140, 100–120, 80–100, 60–80, 40–60, 20–40, and 0–20.




Fig. 6. From top to bottom, varying bands represent the following VoIP bandwidth requirements (4-byte headers): 70–80, 60–70, 50–60, 40–50, 30–40, 20–30, 10–20, 0–10.

A. Delay

Transmission time includes delay due to codec processing as well as propagation delay. ITU-T Recommendation G.114 [8] recommends the following one-way transmission time limits for connections with adequately controlled echo (complying with G.131 [7]):

- 0 to 150 ms: acceptable for most user applications;
- 150 to 400 ms: acceptable for international connections;
- > 400 ms: unacceptable for general network planning purposes; however, it is recognized that in some exceptional cases this limit will be exceeded.

ITU-T Recommendation G.114 Annex B describes the results of subjective tests to evaluate the effects of pure delay on speech quality. A test completed in 1989 showed the percent of users rating the call as poor or worse (POW) for overall quality started increasing above 10% only for delays greater than 500 ms, but POW for interruptability was above 10% for delays of 400 ms. One of the tests, completed in 1990,

“was designed to obtain subjective reactions, in context of interruptability and quality, to echo-free telephone circuits in which various amounts of delay were introduced. The results indicated that long delays did not greatly reduce mean opinion scores over the range of delay tested, viz. 1 to 1000 ms of one-way delay... However, observations during the test and subject interviews after the test showed the subjects experienced some real difficulties in communicating at the longer delays, although subjects did not always associate the difficulty with the delay” [8].

A Japanese study in 1991 measured the effect of delay using six different tasks involving more or less interruptions in the dialogue. The delay detectability threshold was defined as the delay detected by 50% of a task’s subjects. As the interactivity required by the tasks decreased, the delay detectability threshold increased from 45 to 370 ms of one-way

Table 2 Delay Budget for VoIP Using G.729 Codec

Delay Source (G.729)	On-net Budget (ms)
Device Sample Capture	0.1
Encoding Delay (Algorithmic Delay + Processing Delay)	17.5
Packetization/ Depacketization Delay	20
Move to Output Queue/Queue Delay	0.5
Access (up) Link Transmission Delay	10
Backbone Network Transmission Delay	Dnw
Access (down) Link Transmission Delay	10
Input Queue to Application	0.5
Jitter Buffer	60
Decoder Processing Delay	2
Device Playout Delay	0.5
Total	121.1 + Dnw

delay. As the one-way delay increased from 100 to 350 ms, the MOS connection quality decreased from 3.74 (± 0.52) to 3.48 (± 0.48), and the connection acceptability decreased from 80% to 73% [8].

Delay variation, sometimes called jitter, is also important. The receiving gateway or telephone must compensate for delay variation with a jitter buffer, which imposes a delay on early packets and passes late packets with less delay so that the decoded voice streams out of the receiver at a steady rate. Any packets that arrive later than the length of the jitter buffer are discarded. Since we want low packet loss, the jitter buffer delay is the maximum delay variation that we expect. This jitter buffer delay must be included in the total end-to-end delay that the listener experiences during a conversation using packet telephony.

B. Delay Budget

Packetized voice has larger end-to-end delays than a TDM system, making the above delay objectives challenging. A sample on-net delay budget for the G.729 (8 kb/s) codec is shown in Table 2. This budget is not precise. The allocated jitter buffer delay of 60 ms is only an estimate; the actual delay could be larger or smaller. Since the sample budget does not include any specific delays for header compression and decompression, we may consider that, if those functions are employed, the associated processing delay is lumped into the access link delay.

In the absence of Network QoS, the jitter buffer delay could be larger. With QoS and an adaptive jitter buffer, the delay could adapt down to a lower value during a long conversation. This delay budget allows us to stay within the G.114 guidelines, leaving 29 ms for the one-way backbone network delay (Dnw) in a national network. This is achievable in small countries. Network delays in the Asia Pacific region, as well as between North America and Asia, may be higher than 100 ms. According to G.114, these delays are acceptable for international links. However, the end-to-end delays for VoIP calls are considerably larger than for PSTN calls.

IV. NETWORK QOS

There are various approaches to providing QoS in IP networks. Before discussing the QoS options, one must consider whether QoS is really necessary.

Some Internet engineers assert that the way to provide good IP network performance is through provisioning, rather than through complicated QoS protocols. If no link in an IP network is ever more than 30% occupied, even in peak traffic conditions, then the packets should flow through without any queue delays, and elaborate protocols to give priority to one class of packet are not necessary. The design engineer should consider the capacity of the router components to forward small voice packets as well as the bandwidth of the inter-router links in determining the occupancy of the network. If the occupancy is low, then performance should be good. Essentially, the debate is over whether excess network capacity (including link bandwidth and routers) is less expensive than QoS implementation.

The development of QoS features has continued because of the perception of some network engineers that real-time traffic (as well as other applications) may sometimes require priority treatment to achieve good performance. In some parts of the world, bandwidth is at least an order of magnitude more expensive than it is in the United States. In some cases, access links may be expensive and broadband access difficult to obtain, so that QoS may be desirable on the access links even if the core network is lightly loaded. Wireless access links are especially expensive, so QoS is important for wireless mobile IP phone calls. QoS can be achieved by managing router queues and by routing traffic around congested parts of the network. Two key QoS concepts are the IntServ [9] and DiffServ. The IntServ concept is to reserve resources for each flow through the network. RSVP [10] was originally designed to be the reservation protocol. When an application requests a specific QoS for its data stream, RSVP can be used to

deliver the request to each router along the path and to maintain router state to provide the requested service. RSVP transmits two types of Flow Specs conforming to IntServ rules. The traffic specification (Tspec) describes the flow, and the service request specification (Rspec) describes the service requested under the assumption that the flow adheres

to the T_{spec} . Current implementations of IntServ allow a choice of Guaranteed Service or Controlled-Load Service. Guaranteed Service [11] involves traffic policing by a leaky token bucket model to control average traffic. Peak traffic is limited by a peak rate parameter ρ and an interval T so that no more ρT bytes are transmitted in any interval T . The packet size is restricted to be in the range $[m, M]$, so that smaller packets are considered to be of size m and packets larger than M are in violation of the contract. A bandwidth requirement is stated, and enough bandwidth is reserved on each hop to satisfy all the requirements of the flow. (The bandwidth requirement may not be the same on each hop [12].) If each node and hop can accept the service request, the flow should be lossless because the queue size reserved for the flow can be set to the length parameter of the token bucket. This service is designed for interactive real-time applications. To use it effectively, one needs a strict and realistic end-to-end delay budget in addition to bandwidth requirements of the flow.

Controlled-Load Service uses the same T_{spec} as Guaranteed Service. However, an R_{spec} is not defined. Flows using this service should experience the same performance as they would in a lightly loaded “best-effort” network. Controlled-Load Service would be appropriate for call admission control and would prevent the delays and packet losses that make real-time traffic suffer when the network is congested.

There are several reasons for not using IntServ with RSVP for IP telephony. Although IntServ with RSVP would work on a private network for small amounts of traffic, the large number of voice calls that IP telephony service providers carry on their networks would stress an IntServ RSVP system. First, the bandwidth required for voice itself is small, and the RSVP control traffic would be a significant part of the overall traffic. Second, RSVP router code was not designed to support many thousands of simultaneous connections per router.

It should be noted, however, that RSVP is a signaling protocol, and it has been proposed for use in contexts other than IntServ. For example, RSVP-TE is a constraint-based routing protocol for establishing LSPs with associated bandwidth and specified paths in an MPLS network [13]. RSVP has also been proposed as the call admission control mechanism for VoIP in differentiated services networks.

A. Differentiated Services

Since IntServ with RSVP does not scale well to support many thousands of simultaneous connections, the IETF has developed a simpler framework and architecture to support DiffServ [14]. The architecture achieves scalability by aggregating traffic into classifications that are conveyed by means of IP-layer packet marking using the DS field in IPv4 or IPv6 headers. Sophisticated classification, marking, policing, and shaping operations need only be implemented at network boundaries. Service provisioning policies allocate network resources to

traffic streams by marking and conditioning packets as they enter a differentiated services-capable network, in which the packets receive a particular PHB based on the value of the DS field.

The primary goal of differentiated services is to allow different levels of service to be provided for traffic streams on a common network infrastructure. A variety of resource management techniques may be used to achieve this, but the end result will be that some packets will receive different (e.g., better) service than others. This will, for example, allow service providers to offer a real-time service giving priority to the use of bandwidth and router queues, up to the configured amount of capacity allocated to real-time traffic.

Despite the term “differentiated services,” the IETF DiffServ working group undertook to define standards that have more generality than specific services. The reason is that if the IETF were to define new standard services, everyone would have to agree on what constitutes a useful service and every router would have to implement the mechanisms to support it. To deploy that new service, you would have to upgrade the entire Internet. Since a router has only a few functions, it makes more sense to standardize forwarding behavior (“send this packet first” or “drop this packet last”). So the DiffServ working group first defined PHBs, which could be combined with rules to create services.³

An important requirement is scalability, since the IETF intended differentiated services to be deployed in very large networks. To achieve scalability, the DiffServ architecture prescribes treatment for aggregated traffic rather than microflows and forces much of the complexity out of the core of the network into the edge devices, which process lower volumes of traffic and lesser numbers of flows.

The DiffServ architecture is based on a simple model where packets entering a network are classified and possibly conditioned at the boundaries of the network, and then assigned to different behavior aggregates. Each behavior is identified by a single DS codepoint. Within the core of the network, packets are forwarded according to the PHB associated with the DS codepoint.

One candidate PHB for voice service is EF. The objective of the EF PHB is to build a low-loss, low-latency, low-jitter, assured bandwidth, end-to-end service through DS domains. Such a service would appear to endpoints like a point-to-point connection or “virtual leased line.” Since router queues cause traffic to experience loss, jitter, and excessive latency, EF PHB tries to ensure that all EF traffic experiences either no or very small queues. Since queues arise when the short-term traffic arrival rate exceeds the departure rate at some node, this ensures that, at every node, the aggregate EF traffic maximum arrival rate is less than the EF minimum departure rate [15]–[17]. The original idea was to ensure low delay and no packet loss. Subsequent analysis has shown that, under the no loss hypothesis, evaluating the worst-case arrival patterns on each node leads to poor delay bounds after just a few hops. Using a worst-case analysis to

determine admission criteria would lead to unacceptably low utilization.

However, simulations and early EF trials show that good performance can be achieved with reasonable efficiency [18]. The appeal of DiffServ is that it is relatively simple (compared to IntServ), yet provides applications like VoIP some improvement in performance compared to “best-effort” IP networks.

Recently, the IETF DiffServ Working Group has started considering per domain behaviors, but as of this writing the work is still in progress.

However, DiffServ relies on ample network capacity for EF traffic and makes use of standard routing protocols that make no attempt to use the network efficiently. Confronted with network congestion, EF would drop packets at the edge instead of queuing or rerouting them. DiffServ has no topology-aware admission control mechanism. The IETF DiffServ Working Group has not recommended a mechanism for rejecting additional VoIP calls if accepting them would degrade the quality of calls in progress.⁴

B. MPLS-Based QoS

For several decades, traffic engineering and automated rerouting of telephone traffic have increased the efficiency and reliability of the PSTN. Frame relay and ATM also offer source (or “explicit”) routing capabilities that enable traffic engineering. However, IP networks have relied on destination-based routing protocols that send all the packets over the shortest path, without regard to the utilization of the links comprising that path. In some cases, links can be congested by traffic that could be carried on other paths comprised of underutilized links. It is possible to design an IP network to run on top of a frame relay or ATM (“Layer2”) network, providing some traffic engineering features, but this approach adds cost and operational complexity. MPLS offers IP networks the capability to provide traffic engineering as well as a differentiated services approach to voice quality. MPLS separates routing from forwarding, using label swapping as the forwarding mechanism. The physical manifestation of MPLS is the LSR. LSRs perform the routing function in advance by creating LSPs connecting edge routers. The edge router (an LSR) attaches short (four-byte) labels to packets. Each LSR along the LSP swaps the label and passes it along to the next LSR. The last LSR on the LSP removes the label and treats the packet as a normal IP packet.

MPLS LSPs can be established using LDP [19], RSVP-TE [20], or CR-LDP [21]. When using LDP, LSPs have no associated bandwidth. However, when using RSVP-TE or CR-LDP, each LSP can be assigned a bandwidth, and the path can be designated for traffic engineering purposes. MPLS traffic engineering (MPLS-TE) combines extensions to OSPF or IS-IS, to distribute link resource constraints, with the label distribution protocols RSVP-TE or CR-LDP. Resource and policy attributes are configured

on every link and define the capabilities of the network in terms of bandwidth, a Resource Class Affinity string, and a traffic engineering link metric. When performing the constraint-based path computation, the originating LSR compares the link attributes received via OSPF or IS-IS to those configured on the LSP.

Differentiated services can be combined with MPLS to map DiffServ Behavior Aggregates onto LSPs [22]. QoS policies can be designated for particular paths. More specifically, the EXP field of the MPLS label can be set so that each label switch/router in the path knows to give the voice packets highest priority, up to the configured maximum bandwidth for voice on a particular link.

Indeed, the working group co-chairs probably did not believe that admission control was within their charter.

When the high-priority bandwidth is not needed for voice, it can be used for lower priority classes of traffic. DiffServ and MPLS DiffServ are implemented independently of the routing computation. MPLS-TE computes routes for aggregates across all classes and performs admission control over the entire LSP bandwidth. MPLS-TE and MPLS DiffServ can be used at the same time. Alternatively, DiffServ can be combined with traffic engineering to establish separate tunnels for different classes. DS-TE makes MPLS-TE aware of DiffServ, so that one can establish separate LSPs for different classes, taking into account the bandwidth available to each class. So, for example, a separate LSP could be established for voice, and that LSP could be given higher priority than other LSPs, but the amount of voice traffic on a link could be limited to a certain percentage of the total link bandwidth. This capability is currently being standardized by the IETF Traffic Engineering Working Group [23], [24].

Voice DS-TE tunnels can be based on a delay metric or a bandwidth metric. Combining DS-TE with DiffServ over MPLS allows QoS for VoIP with the capability of fast reroute if a link or node failure occurs. DiffServ can guarantee that a specified amount of voice bandwidth is available on each link in a network. DS-TE routing and admission control can create a guaranteed bandwidth tunnel that has the required bandwidth in the highest priority queue on every link. Service conditioning at the edge can ensure that the aggregate VoIP traffic directed onto the guaranteed bandwidth tunnel is less than the capacity of the tunnel. This allows a tight SLA with admission control without overprovisioning the network.

A VoIP network designer can choose DiffServ, MPLS-TE plus DiffServ, or DS-TE according to the economics of the situation. If VoIP is to be a small portion of the total traffic, DiffServ or MPLS-TE plus DiffServ may be sufficient. DS-TE promises more efficient use of an IP network carrying a large proportion of VoIP traffic, with perhaps more operational complexity.

VI. TELEPHONY ROUTING OVER IP (TRIP)

For many years to come, there will be more telephones served by the global PSTN than by IP telephony. Users of IP phones will want to call people who use traditional tele-phones. There are an increasing number of gateways that support VoIP on one side and are connected to the PSTN on the other. Many gateways could complete a call. How does the system find the right gateway?

Telephony routing over IP (TRIP) addresses the following problem: “given a phone number that corresponds to a terminal on a circuit switched network, determine the IP address of a gateway capable of completing a call to that phone number”[34] This is essentially an address to route translation problem.

TRIP does not help find the IP address of a personal computer that serves as an interface to a telephone. For example, a service provider might want to deliver an instant message to a PC associated with a telephone. Directory protocols are better suited to such a problem.

TRIP also does not facilitate calls from a traditional phone to a personal computer that may be used for VoIP. Since IP addresses are often assigned by DHCP or by dialup network access servers, it seems to be a good idea to assign a permanent telephone number to a VoIP terminal, even if that terminal is a computer. A PSTN switch would have to obtain a mapping from this telephone number to an IP address for the PC. This is a name-to-address translation problem that can also be solved using a directory protocol.

The problem that TRIP does address is a complex one. Given the universal connectivity of the PSTN, nearly any VoIP/PSTN gateway could potentially complete a phone call to anywhere in the world. However, there are many factors that influence the decision of which gateway to choose. The calling party may be using signaling or media protocols that are not supported by all gateways. Capacity must also be taken into account in the gateway selection process. Some gateways may support thousands of simultaneous calls, while others support very few. The gateway service provider will want to charge enough to offset costs and make a profit. The user has to pay something, and the gateway service provider has to be paid. However, the end user may be a customer of an IP Telephony service provider who does not own the gateway, but has some business relationship with the gateway service provider. The primary IP telephony service provider may have some gateways as well and is likely to have some policy about what calls are routed to its own gateways and what calls are routed to business partner gateways. Because of these complexities, there cannot be a universal gateway directory. Service providers must exchange information on the availability of gateways, subject to policy. Using this information, each service provider can create its own local database of available gateways.

The main functional component of TRIP is the LS, a logical entity that has access to the telephony routing infor-

mation base (TRIB). The TRIB combines information on gateways available from within its telephony administrative domain with information on gateways available (based on policy) in other IT administrative domains.

TRIP is modeled after the IETF interdomain routing protocol BGP-4 [35], in that it is a protocol for sharing reachability information across administrative domains. As border routers use BGP-4 to distribute IP routes across IP administrative domains, so location servers can use TRIP to distribute telephone routes among telephony administrative domains. “TRIP uses BGP’s interdomain transport mechanism, BGP’s peer communication, BGP’s finite state machine, and similar formats and attributes as BGP” [36] However, TRIP also has some link state features and uses intradomain flooding similar to OSPF. There are some other important differences between BGP and TRIP.

- TRIP is an application layer protocol, whereas BGP is a network layer protocol.
- There may be many intermediate network and IP service providers between location servers that run TRIP. BGP usually runs between routers in adjacent networks.
- TRIP peers exchange information describing routes to application layer location servers.
- TRIP uses a transport network to communicate between servers. It has nothing to do with routing table advertisements.
- There may be islands of TRIP connectivity. There may not be VoIP connectivity among the islands, but within each island, any gateway can have complete connectivity to the entire PSTN.
- Compared to IP routes, many more parameters are necessary to describe gateway routes. Hence gateway routes are relatively more complex.

To illustrate the TRIP architecture, Fig. 18 shows a diagram of the relationship of three ITADs. Each ITAD has at least one LS. ITAD1 has both end users and gateways. ITAD2 has only end users. ITAD3 has only gateways. An LS learns about the gateways in their domain through an out-of-band intradomain protocol, which is represented by the dashed lines in ITAD3. The administrative domains have agreements that allow the LSs to exchange gateway data. Using TRIP, the LS in ITAD2 can learn about the three gateways in ITAD3, as well as the two gateways in ITAD1. The end users in ITAD2 can use a non-TRIP protocol to access the LS databases. The LS in ITAD1 can learn about the gateways in ITAD3 from the LS in ITAD2; this information might be in an aggregated advertisement.

A. Example — Clearinghouse

A clearinghouse is like a route reflector. Members of the clearinghouse agree to accept each other’s IP telephony traffic at their gateways. Clearinghouse members can use

TRIP to exchange routes with the clearinghouse. Fig. 19




Fig. 18. TRIP architecture.

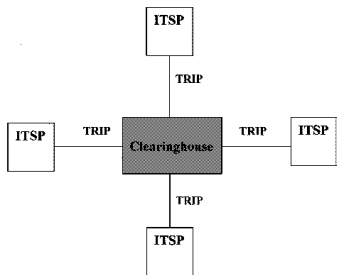


Fig. 19. IT clearinghouse using TRIP.

shows a diagram of four ITSPs using TRIP to exchange gateway routes with the clearinghouse.

VII. VOIP ISSUES WITH NAT AND FIREWALLS

VoIP is one of many IP applications that have problems traversing NATs and firewalls. While there are solutions, they all increase the expense and operational complexity of Internet telephony.

NAT allows private networks to connect to a common network (e.g., the Internet) although they have overlapping address realms. NAT is used and tolerated as a means to ameliorate IPv4 address depletion by allowing globally registered IP addresses to be reused or shared by several hosts. NAT also protects the privacy of the internal network topology and addresses. NAT routers, placed at the border between private and public networks, convert the private addresses in each IP packet into IANA-registered public IP addresses. In addition to modifying the IP address, NAT must modify the IP checksum and the TCP checksum. The packet sender and receiver should remain unaware that NAT is taking place. Firewalls commonly support NAT.

There are both static and dynamic NAT devices and routers, but dynamic NAT is more common today. Edge devices that run dynamic NAT allow an entire private IP subnet to share a pool of public IP addresses. So long as a private host has an outgoing connection, incoming packets sent to the public NAT address can reach it. After the connection is terminated or times out, the binding expires, and the NAT returns the address to the pool for reuse. Network Address Port Translation (NAPT), a variation of dynamic NAT, allows many hosts to share a single IP address by multiplexing streams differentiated by

TCP/UDP port number. For example, suppose private hosts 10.0.0.2 and 10.0.0.3 both send packets from source port 1180. A NAPT router might translate these to a single public IP address 9.245.160.1 and two different source ports, say 5431 and 5432. The NAPT would route response traffic for port 5431 to 10.0.0.2:1180, while traffic to port 5432 would go to 10.0.0.3:1180.

Multihost residential users, teleworkers, and small businesses use NAPT devices (sometimes called SOHO routers) to allow multiple computers to share a single public IP address for outbound traffic while blocking inbound session requests. A provider of DSL or cable modem service often assigns the single IP address. A NAPT router allows several computers to share that IP address. Enterprises with private address realms also use NAPT

A. Protocol Complications With NAT

VoIP is one of many applications that can be adversely affected when IP clients connect through a NAT or NAPT. The NAT device may use an application level gateway (ALG). An ALG examines and modifies application payload content to allow packets from a specific application or protocol to pass through the NAT transparently. However, few NAT devices offer ALG functions for VoIP, and some protocols are not amenable to this approach.

There are several categories of problems that VoIP applications have with NAT.

- 1) Many applications fail with NAT because the packets contain IP address or port information in the payload. A simple NAT only changes the IP address of the packet itself, not the IP addresses and ports in the payload. In the case of H.323, it is the call setup packets that contain the address and port information in the payload.
- 2) H.323 and SIP, as well as other applications such as FTP and RTSP, use bundled sessions. They exchange address and port parameters within a control session to establish data sessions. NAT cannot determine the inter-dependency of the bundled sessions and assigns unrelated addresses and port numbers to these sessions, which does not work.
- 3) An IP application (such as IP phone) that attempts to originate a session from an external realm will be able to locate its peer in a private realm only when it knows the externally assigned IP address ahead of time. This is a problem for a traditional dynamic NAT, which only permits sessions to be established in one direction.
- 4) SIP messages may carry URL's that specify signaling addresses in the "Contact," "To," and "From" fields. Once they traverse a NAT, the IP addresses and domain names in the host port portion of the URL may not be valid.

B. H.323 Characteristics

H.323 is a protocol suite that uses multiple UDP streams and dynamic ports. An H.323 call consists of many different

simultaneous connections. There are two or more TCP connections for each call. For a voice conference call, there may be as many as four different UDP ports open. All connections except one are made to dynamic ports.

During call setup, a TCP connection carries H.225 signaling, including the Q.931 messages. During slow start call setup, the H.245 messages carry the terminal characteristics and requested call parameters in a TCP connection separate from the H.225 data stream. There is no well-known port associated with the H.245 channel. Instead, the H.225 channel is used to convey the H.245 port information. The firewall needs to monitor the H.225 channel for the H.245 port, because it is not possible to implement a sufficiently stringent static rule that allows an H.245 connection while blocking other undesired TCP connections. During FastStart call setup, the H.245 message is imbedded in the H.225 message along with the Q.931 message. To work properly, an ALG has to modify the addresses inside these messages. Q.931 and H.245 messages are encoded in ASN.1 in the packet payload, and they are variable in length. Of course, these difficulties have not prevented vendors from developing NAT-enabled firewalls with ALG functions that allow H.323 to pass through. However, small inexpensive NATs and firewalls do not have H.323 ALGs.

C. NAT/Firewall Problems With RTP

Media transport for all IP multimedia applications, including VoIP, uses RTP in conjunction with UDP. There are no fixed ports associated with RTP, and it is impossible to define static rules that can allow RTP media through a firewall without also allowing undesirable packets to pass through. Furthermore, RTP and RTCP ports are paired, with RTP receiving an even port number, and RTCP receiving the next higher odd port number. NAPT typically assigns new port numbers at random, breaking the pair relationship of RTP and RTCP port numbers. Also, for multimedia sessions, the NAT functions scramble the source and destination addresses used for packets and without special processing by the NAT, these will not correspond with the values used in the control connections. Thus, the multimedia devices may not associate the RTP sessions with the correct call.

D. NAT/Firewall Traversal

We have observed some problems that session-oriented protocols such as VoIP experience with NATs and firewalls. There are four types of solutions.

The first solution is a proxy placed at the border between two domains (e.g., between a private IP address space and a public address space). The proxy would terminate sessions with both hosts, or with both client and server, and relay application signaling messages as well RTP media streams transparently between the two hosts. Only designated

protocols, such as SIP or H.323, would pass through the proxy. All other traffic would have to traverse the NAT and/or firewall to communicate between the two domains.

The second solution is an ALG embedded in the NAT or firewall. The ALG does not terminate sessions, but rather examines and modifies application payload content to allow VoIP traffic traverse the NAT/firewall. The ALG is the most common commercial solution now, but ALG-enabled firewalls tend to be somewhat expensive. Placing several ALGs within the same firewall increases its complexity and may degrade performance. Furthermore, any changes in the VoIP protocol used will require a new ALG from the firewall vendor for all the previously installed firewalls that VoIP has to traverse. The upgrade also tends to be expensive.

A third approach is to remove the application logic from the NAT/firewall. A new type of firewall dynamically opens "pinholes" to let a VoIP call through it, without exposing the private network by allowing penetration by a wide range of IP addresses. A firewall control proxy (FCP), placed in the signaling path between private and public domains, monitors the call setup signals (such as H.323 and SIP) and commands the firewall to allow RTP streams destined to the appropriate IP addresses to pass through. For protocols such as SIP and H.323, moving stateful inspection and manipulation of signaling packets out of NAT/firewalls should improve scalability and performance while reducing development costs. The IETF is exploring this third approach in the Middlebox Communications (Midcom) Working Group. The MidCom group is trying to agree on a control protocol that would enable another device (an FCP, basically) to control middle boxes such as NATs and firewalls. By providing a generalized standard interface communications interface for the middle boxes, the working group hopes to improve performance, lower software development and maintenance costs, and easier deployment of new applications. [37]

These two types of solutions, ALGs and FCP/MidCom, require changes to NAT and firewall design. A fourth type of solution seeks a means to "traverse" the NAT and/or firewall without changing its design, and without requiring it to perform additional processing. The challenge of this type of solution is to allow VoIP signaling and media streams to traverse the NAT and/or firewall without compromising security.

Two Internet drafts [38], [39] have suggested ways to allow VoIP and other multimedia traffic to traverse NATs and firewalls. Although the methods are different, they both employ external proxy servers with persistent connections to the VoIP/multimedia devices. Two essential elements of these traversal methods are as follows.

- 1) The user behind the NAT must send the first packet to establish the NAT binding.
- 2) Media sent to user A must be to the source port from which A's media came.

To that end, devices in the private address realms communicate with the proxy servers in the public address realm via "probe packets" or "cookies." The proxy servers

associate the origination address/port pair with the “token” or “cookie.”

SUMMARY AND CONCLUSION

Providing reliable, high-quality voice communications over a network designed for data communications is a complex engineering challenge. Factors involved in designing a high-quality VoIP system include the choice of codec and call signaling protocol. There are engineering tradeoffs between delay and efficiency of bandwidth utilization. Packetized voice has larger end-to-end delays than a TDM system. One reason is that an IP network typically has higher delay variation than a TDM system. Since any packets that arrive later than the length of the jitter buffer are discarded, the jitter buffer delay must be set to the maximum delay variation that we expect, in order to achieve low packet loss probability. The jitter buffer delay becomes a major component of the end-to-end delay budget, to which must be added the encoding delay and packetization delay. VoIP performance can be improved by network QoS techniques (such as differentiated services) that are not widely available in the public Internet today, but may be deployed by specialized commercial IP networks.

We have reviewed the motivation and characteristics of TRIP, a location server protocol for the inter-domain advertising of PSTN destinations reachable from participating gateways, and the attributes of those gateways. We also reviewed the challenges that VoIP signaling protocols and media packet streams have in coping with network address translation and firewalls.

While posing complex engineering challenges, VoIP remains a topic of extensive product development and intense standards activity. We can expect more VoIP solutions and more protocol developments in the near future, as well as an increasing volume of telephone traffic using this technology.

REFERENCES

[1] M. Perkins, K. Evans, D. Pascal, and L. Thorpe, “Characterizing the subjective performance of the ITU-T 8 kb/s speech coding algorithm – ITU-T G.729,” *IEEE Commun. Mag.*, vol. 35, pp. 74–81, Sept. 1997.

[2] H. Schulzrinne, S. Casner, R. Frederick, and V. Jacobson, “RTP: A transport protocol for real-time applications,” *IETF RFC 1889*, 1996.

[3] A. Benyassine, E. Schlotmot, H. Y. Su, D. Massaloux, C. Lamblin, and J. P. Petit, “ITU-T G.729 annex B: A silence compression scheme for use with G.729 optimized for V.70 digital simultaneous voice and data applications,” *IEEE Commun. Mag.*, vol. 35, pp. 64–73, Sept. 1997.

[4] M. Degermark, B. Nordgren, and S. Pink, “IP Header Compression,” *IETF RFC 2507*, 1999.

[5] S. Casner and V. Jacobson, “Compressing IP/UDP/RTP headers for low-speed serial links,” *IETF RFC 2508*, 1999.

[6] M. Engan, S. Casner, and C. Bormann, “IP header compression over PPP,” *IETF RFC 2509*, 1999.

[7] “Stability and Echo,” *CCITT Recommendation G.131*, 1988.

[8] “One-way transmission time,” *ITU-T Recommendation G.114*, 1996.

[9] R. Braden, D. Clark, and S. Shenker, “Integrated services in the internet architecture: An overview,” *IETF RFC 1633*, 1994.

[10] R. Braden, L. Zhang, S. Berson, S. Herzog, and S. Jamin, “Resource reservation protocol (RSVP) version 1 functional specification,” *IETF RRC 2205*, 1997.

[11] S. Shenker, C. Partridge, and R. Guerin, “Specification of guaranteed quality of service,” *IETF RFC 2212*, 1997.

[12] P. White and J. Crowcroft, “The integrated services in the internet: State of the art,” *Proc. IEEE*, vol. 85, pp. 1934–1946, Dec. 1997.

[13] D. Awduche, A. Hannan, and X. Xiao, “Applicability statement for extensions to RSVP for LSP tunnels,” *IETF RFC 3210*, 2001.

[14] D. Black, S. Blake, M. Carlson, E. Davies, Z. Wong, and W. Weiss, “An architecture for differentiated services,” *IETF RFC 2475*, 1998.

[15] V. Jacobson, K. Nichols, and K. Poduri, “An expedited forwarding PHB,” *IETF RFC 2598*, 1999.

[16] B. Davie and A. Charney et al., “An expedited forwarding PHB,” *IETF RFC 3246*, 2002, to be published