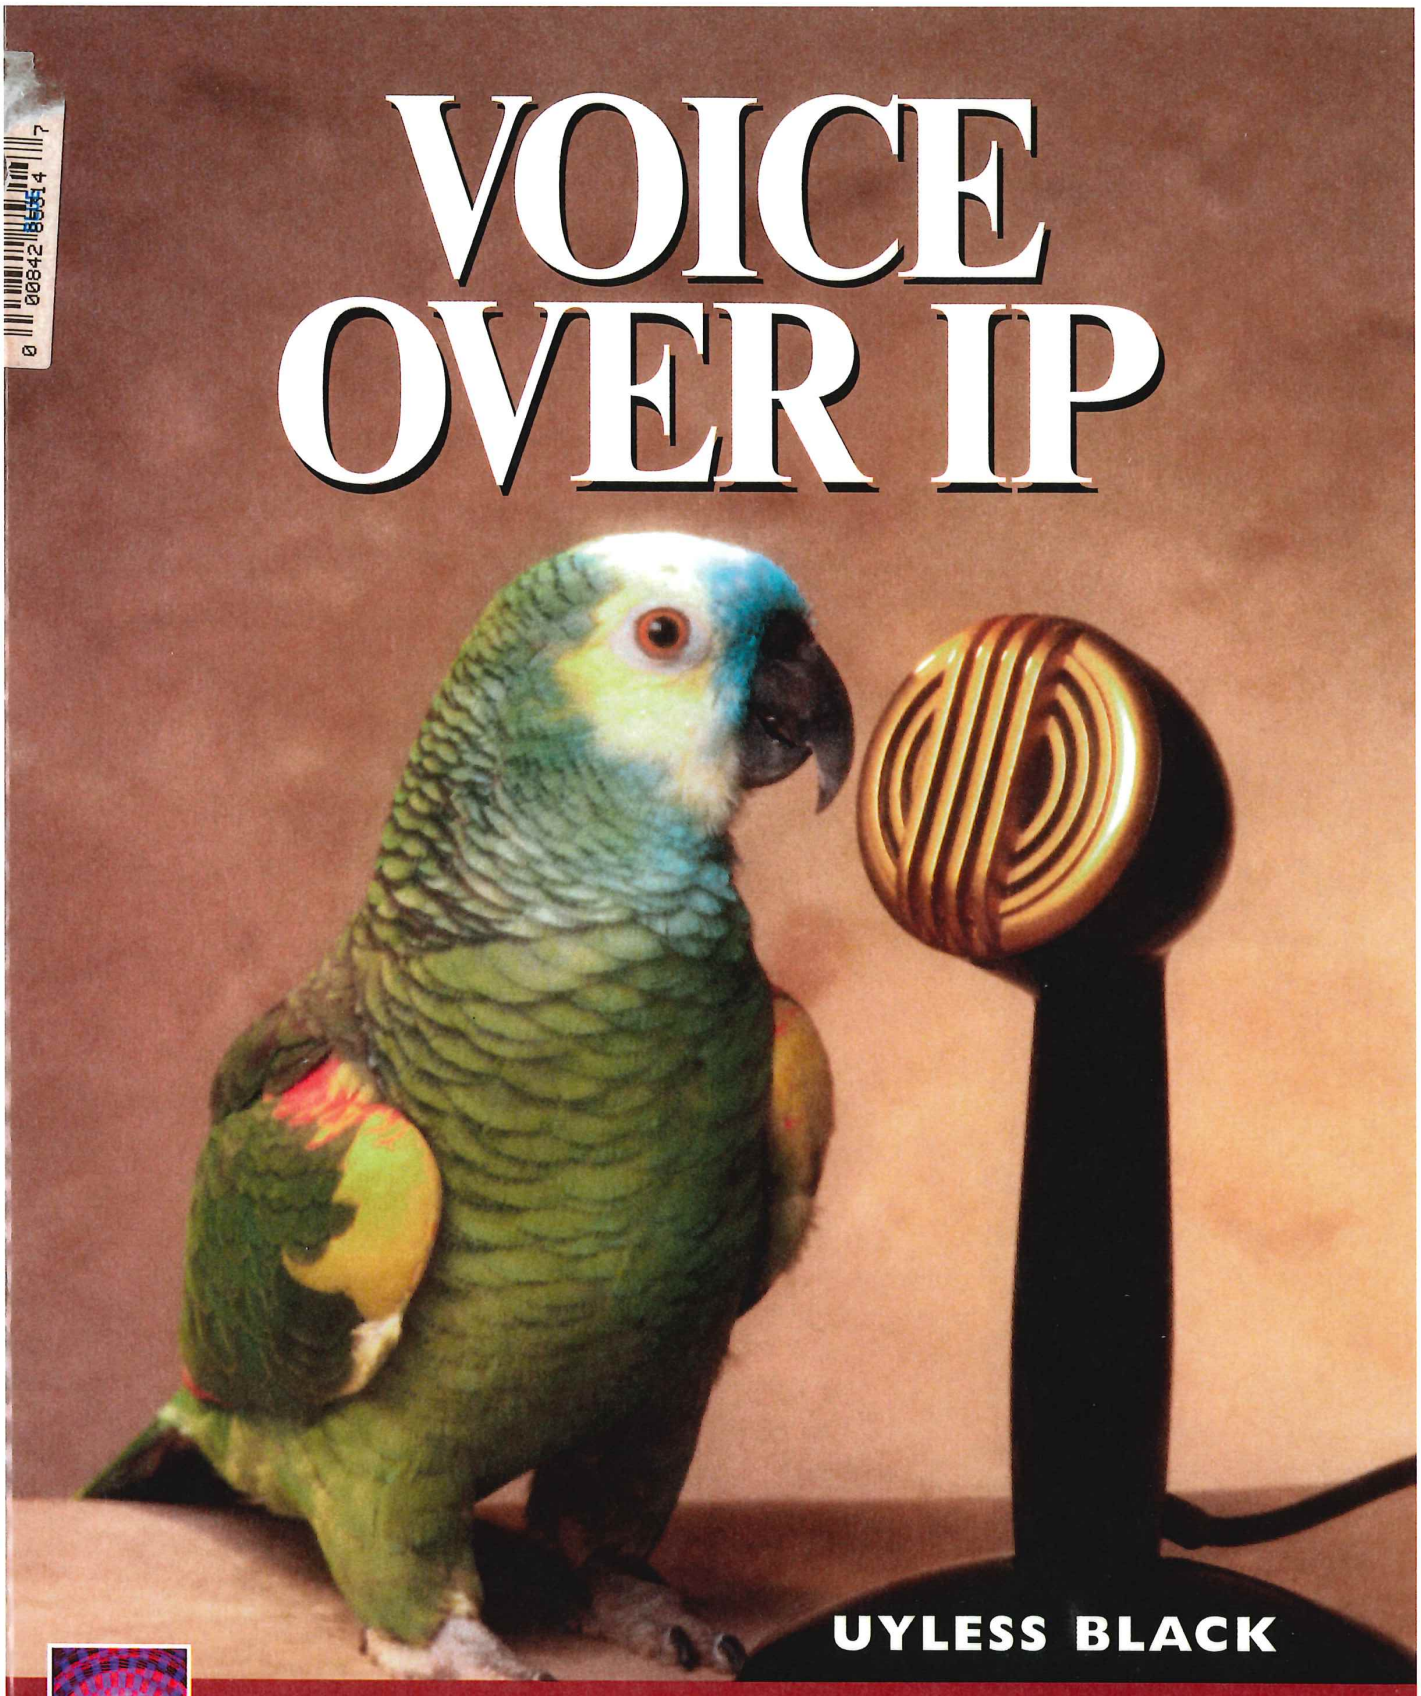
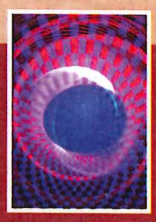
VOICE OVER IP

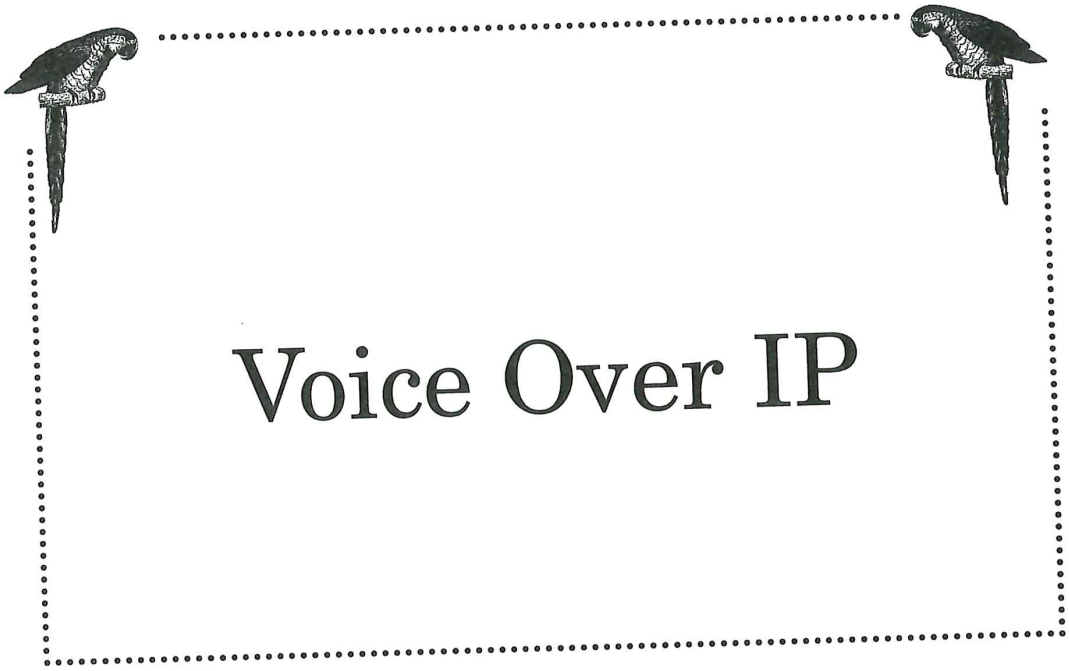


UYLESS BLACK



PRENTICE HALL SERIES
IN ADVANCED COMMUNICATIONS TECHNOLOGIES

Petitioner's Exhibit 1021
Cisco Systems, Inc. v. Video Solutions Pte. Ltd
Page 000001



Voice Over IP

UYLESS BLACK

Prentice Hall PTR
Upper Saddle River, New Jersey 07458
http://www.phptr.com

Petitioner's Exhibit 1021
Cisco Systems, Inc. v. Video Solutions Pte. Ltd
Page 000002

Library of Congress Cataloging-in-Publication Data

Black, Uyless D.
Voice over IP / Uyless Black ; editor, Mary Franz.
p. cm.
Includes bibliographical references.
ISBN 0–13–022463–4
1. Internet telephony. 2. TCP/IP (Computer network protocol)
I. Title.
TK5105.8865.B53 1999
004.6—dc21 99–31326
CIP

Acquisitions editor: Mary Franz
Editorial assistant: Noreen Regine
Cover designer: Talar Agasyan
Cover design director: Jerry Votta
Manufacturing manager: Alexis R. Heydt
Marketing manager: Lisa Konzelmann
Project coordinator: Anne Trowbridge
Compositor/Production services: Pine Tree Composition, Inc.

© 2000 by Prentice Hall PTR
Prentice-Hall, Inc.
Upper Saddle River, New Jersey 07458

Prentice Hall books are widely used by corporations and government agencies for training, marketing, and resale.

The publisher offers discounts on this book when ordered in bulk quantities. For more information contact:

Corporate Sales Department
Phone: 800–382–3419
Fax: 201–236–7141
E-mail: corpsales@prenhall.com

Or write:

Prentice Hall PTR
Corp. Sales Dept.
One Lake Street
Upper Saddle River, New Jersey 07458

All rights reserved. No part of this book may be reproduced, in any form or by any means, without permission in writing from the publisher.

Printed in the United States of America
10 9 8 7 6 5 4 3

ISBN: 0–13–022463–4

Prentice-Hall International (UK) Limited, *London*
Prentice-Hall of Australia Pty. Limited, *Sydney*
Prentice-Hall Canada Inc., *Toronto*
Prentice-Hall Hispanoamericana, S.A., *Mexico*
Prentice-Hall of India Private Limited, *New Delhi*
Prentice-Hall of Japan, Inc., *Tokyo*
Prentice-Hall (Singapore) Pte. Ltd., *Singapore*
Editora Prentice-Hall do Brasil, Ltda., *Rio de Janeiro*

- VoIP over Fast or Gigabit Ethernet
- VoIP over AAL 1/AAL 2/AAL 5 over ATM over SONET
- VoIP over Frame Relay
- VoIP over FDDI
- VoIP over RTP over UDP, then over IP, and layers 2 and 1

We could go on, for there are other choices and combinations. The purpose of this discussion is to emphasize that different bearer and teleservices products will be available to meet diverse customer requirements. Later chapters will provide overviews of some of these technologies, but the main emphasis in this book is VoIP itself. Other books in this series are dedicated to the detailed discussion of these supporting technologies.

BASIC TERMS AND CONCEPTS

Some of the readers of this book are familiar with the Internet, and others are not. I use the next part of this chapter to provide a tutorial on some basic Internet terms and concepts. The experienced reader can skip to the section titled, Evaluation the Factors in Packetized Voice (page 25).

The Internet is an association of thousands of user computers that communicate with each other through networks. These user computers are called "hosts." The networks are connected together through another machine that relays the host computer traffic between user applications (such as e-mail and file transfer) that are running on the hosts. The Internet uses the term *Gateway* or *router* to describe the machine that performs the relaying functions between networks. Figure 1–9 shows a Gateway/router placed between networks A, B, and C. Routers A, B, and C are said to be internetworking machines, since they connect networks together.

Networks are often grouped together, and the individual networks are called subnetworks. They are full networks unto themselves, but the idea allows a set of subnetworks (subnets) to be associated with one organization, or some type of administrative domain, such as an Internet Service Provider (ISP). The organization can identify each network with a subnet identifier (ID) and can group these IDs (networks) together (or treat them separately). The grouping concept is called address aggregation. In Figure 1–9, Subnets A.1 and A.2 can be collectively identified with address A.

Petitioner's Exhibit 1021
Cisco Systems, Inc. v. Video Solutions Pte. Ltd
Page 000004

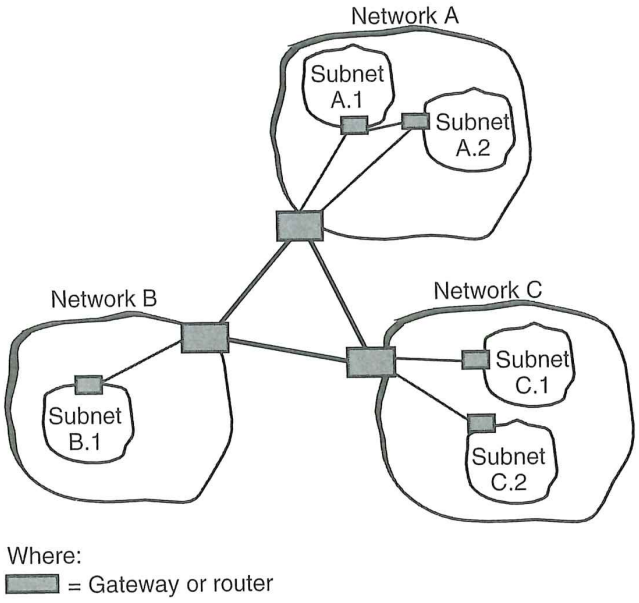


Figure 1–9 Internetworking and Internets

This approach is useful because, like the telephone system, it allows the Internet components to be identified with a hierarchical address. For example, in a telephone system, a person can be reached by dialing first an area code, then an exchange number, and then a subscriber number. In the Internet, addresses are managed by a form of hierarchical aggregation called address prefixes.[10]

An internetworking router is designed to remain transparent to the end-user application. Since the end-user application resides in the host computer, the router need not burden itself with application protocols, and can dedicate itself to fewer tasks, such as managing the traffic between networks.

ATTRIBUTES OF THE INTERNET

The Internet was developed to support the transfer of data traffic (packets) between computers and workstations with the use of adaptive routing features, see Table 1–2. Adaptive routing means the traffic may

[10]The concept of a network address in the Internet is based on the use of class-based IP addresses. With prefixing, address classes are no longer important. For more information see [BLAC98] Black, Uyless. *Advanced Features of the Internet*, Prentice Hall, 1998.

Table 1–2 Attributes of the Internet

Attribute	Consequence(s)
Data applications	Not "tuned" for voice, or video
Adaptive routing:	Path may vary during traffic transfer, and packets may arrive out-of-order
Connectionless:	Circuits are not set up between users
"Best effort" delivery service:	Traffic discarded if problems occur

take different routes through the Internet depending on network conditions at a specific time, such as congestion, a failed link, etc. The possible result of adaptive routing is that the destination user may receive the packets out-of-order. An Internet protocol (the Transmission Control Protocol, TCP) at the receiver can be used to reorder the packets into the proper sequence. The other possible result of adaptive routing is that the arrival rate of the packets at the receiver may vary; some packets may arrive with little delay while others take longer.

The Internet is designed as a connectionless system. This means that there are no "affiliations" established between the machines in the Internet. As a result, the Internet does not maintain ongoing knowledge of the user's traffic, and does not build a fixed path between the switches from the source and to destination host machines. In effect, Internet Protocol (IP) traffic routing is stateless; that is to say, it does not build tables to maintain information about a connection, because there is no connection.

The connectionless aspect of the Internet goes hand-in-hand with the adaptive routing concept. But in the telephone network, the opposite architecture is employed: connection-oriented fixed paths between the calling and called parties. The telephony approach is needed to support the real time, non-varying delay requirements of speech. Whereas, the Internet is a data network, and most data applications do not require the real-time transport service.

The Internet is a "best effort" delivery network. The term best effort means that the Internet will attempt to deliver the traffic, but if problems occur (damaged bits due to noise, congestion at a router, etc.), or the destination host cannot be found, the traffic is discarded. In most instances, TCP in the originating host machine can resend the lost or damaged packets.

Internet Attributes vis-à-vis Voice Traffic

It is evident from this brief examination of the Internet, and our previous discussions about the requirements for toll-quality voice, that

Petitioner's Exhibit 1021
Cisco Systems, Inc. v. Video Solutions Pte. Ltd
Page 000006

the Internet is not a particularly good choice for the transport of voice traffic. However, this is a moot point, since it is the anointed technology for packetized voice. And as stated before, the Internet will evolve to support toll quality voice.

THE INTERNET LAYERED ARCHITECTURE

Many of the concepts in this book are explained with the layered protocol concept. This section provides a brief review of the Internet layers and Chapter 2 gives more detailed information.

Figure 1–10 a reviews the Internet protocol suite layers. With some exceptions, the Open Systems Interconnection (OSI) Model layer 6 is not used. Layer 5 is not used at all.

The physical and data link layers are (as a general rule) also not defined. The philosophy is to rely on existing physical and data link systems. One notable exception to this practice is at the data link layer, where the Internet task forces have defined the Point-to-Point Protocol (PPP).

For the newcomer, here is a summary of the functions of the layers:

- *Physical layer:* Defines the media, and physical aspects of the signals (voltages, etc.). Defines clocking and synchronization operations. Defines physical connectors. Also identified as layer 1 or L_1. Examples are T1, E1, SONET, and L_1 of Ethernet.

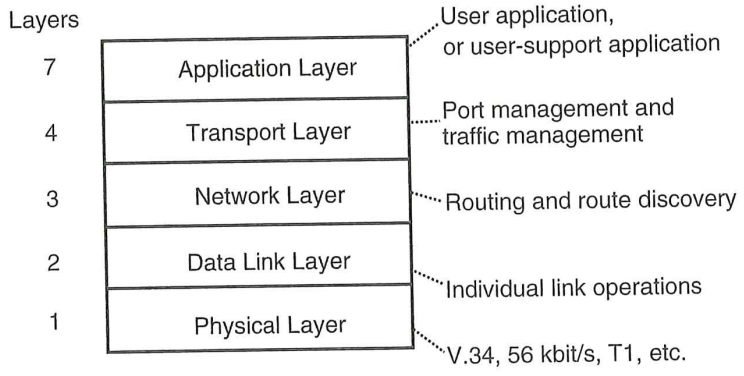


Figure 1–10 The Internet Protocol suite layers

- *Data link layer:* Supports the transfer of traffic over one link. May perform error detection and retransmission, depending on the specific link layer protocol. Also identified as layer 2 or L_2. Examples are PPP, LAPD, and L_2 of Ethernet.
- *Network layer:* Performs forwarding operations and route discovery. Supports some limited diagnostic functions, such as status reports. Also identified as network level, layer 3, or L_3. An example of forwarding is IP. An example of route discovery is Open Shortest Path First (OSPF).
- *Transport layer:* Supports end-to-end acknowledgment of traffic, as an option. At a receiving host, supports the identification (with a port number) of the layer_7 protocol to be invoked to support incoming traffic. Also identified as layer 4 or L_4. Examples are TCP and UDP.
- *Application layer:* Contains the end user application, or another application that directly supports the end user application, such as a file transfer or an email operation. Also identified as layer 7 or L_7.

EVALUATING THE FACTORS IN PACKETIZED VOICE

Now that we have a general understanding of the Internet attributes, as well as the Internet protocol suite, it is appropriate to evaluate how to place speech traffic on an internet. The VoIP designer must evaluate three key factors: (a) packet delay, (b) bandwidth requirements, and (c) computational effort.

Packet delay describes how long it takes to send the packet from the sender to the receiver. Two aspects of packet delay should be known. The first aspect is how long it takes to send the traffic from the sender to the receiver. The second aspect is the variation in time of the arrival of the packets at the receiver. The variation in delay is called jitter.

The second factor deals with how much bandwidth is required to support the voice or video transmission. The bandwidth calculation must factor in the bits required to represent the speech signal as well as the overhead headers (protocol control information) that are used to support the signals. At a minimum this includes the L_2 header, the IP header, the UDP header, the L_7 header, and the headers created by the voice coder. All totaled, they add significant overhead to the voice packet, and

Petitioner's Exhibit 1021
Cisco Systems, Inc. v. Video Solutions Pte. Ltd
Page 000008

of excessive delays in the network, the packets may be discarded because they are of no use if they arrive at the receiver too late. Again, the loss does not severely affect voice fidelity if the lost packets are less than 10 percent of the total packets transmitted. As discussed before, data packets can ill-afford to be lost or discarded.

Tolerance for Delay

Yet another difference between voice and data transmissions deals with network delay. For packetized voice to be translated back to an analog signal in a real-time mode, the two-way delay for voice packets must be constant and generally must be low—usually less than 300 ms. Why is low delay important? If it takes a long time for the voice packets to be sent from the speaker (person A) to the listener (person B), when speaker A stops talking, receiver B is still receiving the speech packets. Person B cannot start talking until all the speech signals have arrived. In the meantime, person A hears nothing for a while.

The two-way delay measures how long it takes: (a) for A's speech to reach B, (b) for B to hear the speech, (c) for B to talk back, and (d) for A to hear B's response. If the delay becomes long (say, over 400 or 500 ms), the conversation appears phony, almost like a half-duplex connection where the two people are taking turns talking, but waiting a while before taking the turn. All-in-all, it can be quite annoying.

For data packets, the network delay can vary considerably. Indeed, the packets can be transmitted asynchronously through the network, without regard to timing arrangements between the sender and the receiver.

Tolerance for Delay and the Effects on Queues. Voice packets require a short queue length at the network nodes in order to reduce delay, or at least to make the delay more predictable. The short voice packet queue lengths can experience overflow occasionally, with resulting packet loss. However, data packets require longer queue lengths to be longer to prevent packet loss in overflow conditions.

Tolerance for Variable Bit Rates (VBRs) and Constant Bit Rates (CBRs)

A useful method to describe the nature of applications traffic is through two concepts known as variable bit rate (VBR) and constant bit rate (CBR). An application using VBR schemes does not require a constant and continuous allocation of bandwidth. These applications are said to be bursty, which means that they transmit and receive traffic asynchro-

Petitioner's Exhibit 1021
Cisco Systems, Inc. v. Video Solutions Pte. Ltd
Page 000009





2

Characteristics of the Internet and IP

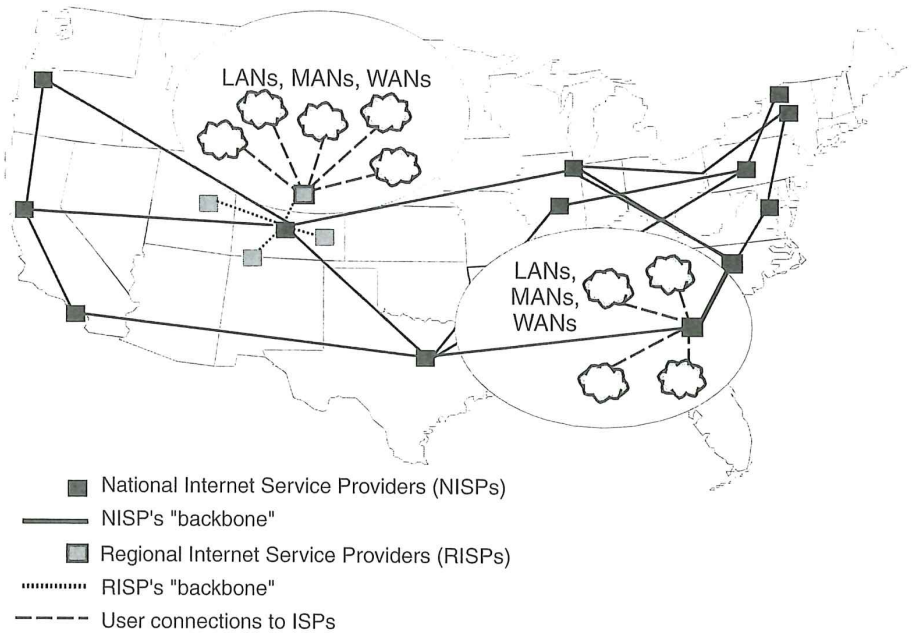
This chapter is divided into three sections. The first section explains the Internet architecture and the characteristics of the Internet; that is, how it "behaves" (or misbehaves) in the handling of user traffic. This section is a more detailed explanation than the generalized discussion in Chapter 1. The next section provides a brief overview of IP. The last section provides a brief overview of Internet ports and sockets and the major features of Transmission Control Protocol (TCP) and User Datagram Protocol (UDP).

The reader who is experienced with the Internet can skip the last two sections, but I recommend all read this first section.

ARCHITECTURE OF AN INTERNET

Today, the Internet is a complex collage of regional and national networks that are interconnected together with routers. The communications links used by the ISPs are leased lines from the telephone system, usually DS1 or DS3 lines, and increasingly SONET lines. Other lines are provided by competitive access providers (CAPs), and still others are provided by private carriers, such as private microwave, and satellite operators.

Petitioner's Exhibit 1021
Cisco Systems, Inc. v. Video Solutions Pte. Ltd
Page 000010



Where:

ISPs	Internet service providers
LANs	Local area networks
MANs	Metropolitan area networks
WANs	Wide area networks

Figure 2–1 Internet overview

The Internet connections and the many ISPs are impossible to depict in one illustration, but Figure 2–1 provides an accurate view of the basic structure of the Internet topology.

ISPs AND THE TELEPHONE NETWORK

The Internet Service Providers (ISPs) provide the access into the Internet. They provide this access through the telephone network, and the Internet user (say from the home) uses a dial-up modem to connect to the ISP through the telco local exchange carrier (LEC), and the LEC central office (CO), see Figure 2–2.

The telephone companies are required to provide the ISPs with connections to the telephone company's customers. The connections are provided at the telco's plant. These connections from the customer to the CO's main distribution frame (MDF, a physical "patch panel") may be

Petitioner's Exhibit 1021
Cisco Systems, Inc. v. Video Solutions Pte. Ltd
Page 000011

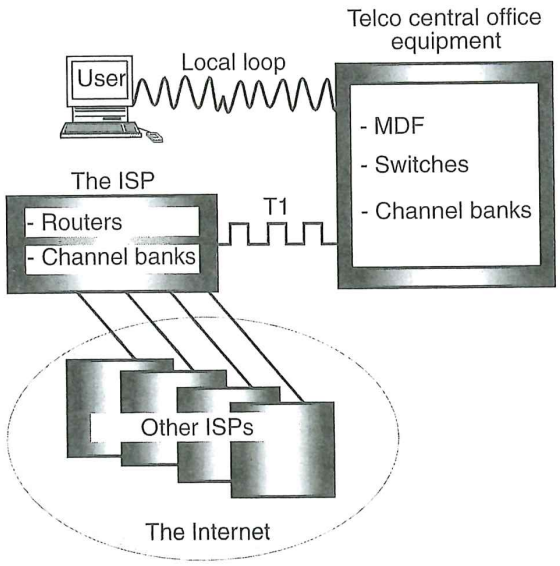


Figure 2–2 Typical telco and Internet service provider (ISP) set-up

patched to the local switch (a conventional circuit switch, for example). Alternately, the line from the customer may be patched to a digital cross connect (DCS) machine, which is a software-based component that can support leased lines—in this example, leased lines to the ISP.

The basic idea of Figure 2–2 is to emphasize that the Internet user gains access to the Internet and the user's ISP through the telco's local loop and CO's equipment. At the CO, the connection to the ISP is attained through the DCS.

The trunk connections from the CO to the ISP are at an agreed upon point of presence (POP), say at the CO's site. On the trunk side, the signals are now either of the T1 family or the SONET STS/OC family (Synchronous Optical Network, Synchronous Transport Signal, Optical Carrier).

ATTRIBUTES OF THE INTERNET

Before we begin analyzing the Internet operations, it will prove helpful to examine in more detail the major attributes of the Internet, introduced in Chapter 1.

As we learned, the Internet was developed to transfer data traffic by using adaptive routing features. By adaptive routing, we mean that the traffic may take different routes through the Internet depending on the

Petitioner's Exhibit 1021
Cisco Systems, Inc. v. Video Solutions Pte. Ltd
Page 000012

conditions at a specific time. In addition, the Internet is designed as a connectionless system which means that there are no "affiliations" established between the machines in the Internet. As a result, the Internet does not maintain an ongoing knowledge of the user's traffic. In effect, the Internet Protocol (IP) is stateless; that is to say, it does not build tables to maintain information about a connection because there is no connection.

The Internet is a "best effort" delivery network. The term best effort means that the Internet will attempt to deliver the traffic, but if problems occur, the traffic is discarded.

Finally, the Internet supports either unicasting (one-to-one) or multicasting (one-to-many) operations. The multicasting feature has proved to be a very useful tool for conference calling as well as for the downline loading of software or data to multiple sites.

Round-Trip Time (RTT)

Round-trip time (RTT) is a measure of the time it takes to send a packet to a destination node and receive a reply from that node. RTT includes the transmission time in both directions and the processing time at the destination node.

Most RTTs in the Internet are within the range of 70–160 ms, although large variations to RTT do occur. Due to the asynchronous nature of the Internet, RTT is not consistent. During periods where Internet traffic is heavy, the RTT may exceed 300 ms.

The ITU-T G.114 recommendation limits RTT to 300 ms or less for telephone traffic. This performance factor is based on many studies and observations; they conclude that longer delays in a telephone-based conversation gives the impression to the conversationalists that they are using a half-duplex circuit. (Interestingly, other surveys show some people tolerate large RTTs of up to 800 ms. But this tolerant population is in the minority.)

PACKET LOSS

Another Internet characteristic that is important to voice and video applications is packet loss. Two factors are involved: (a) how often packet loss occurs and (b) how many successive (contiguous) packets are affected. Packet loss is masked in a data application by using TCP to resend the lost TCP segments. Certainly, the loss of many segments and

Petitioner's Exhibit 1021
Cisco Systems, Inc. v. Video Solutions Pte. Ltd
Page 000013

their retransmissions will affect the application's performance, but on the whole, the end-user data application is not concerned (or aware) of packet loss.

Packet loss is quite important in voice and video applications, since the loss may affect the outcome of the decoding process at the receiver, and may also be detected by the end-user's ears or eyes. Notwithstanding, today's voice coders can produce high-quality voice signals in the face of about 10 percent loss of the voice packets (G.723.1 as the example, see Chapter 4), *if* the packet losses are random and independent. G.723.1 compensates for this loss by using the previous packet to simulate the characteristics of the vocal signal that was in the lost packet.

Traffic loss in the Internet is bursty: large packet losses occur in a small number of bursts. This characteristic of internet behavior complicates the support of telephony, because packetized voice works best if the packet loss is random and independent.

The effect of packet loss can be alleviated somewhat by the use of forward error correction (FEC) schemes, and methods have been devised to compensate for loss bursts [BORE97].[1] These schemes add extra delay to the process, and may result in the loss of the packet because it is made available to the user in a time domain that is too late to be useful.

One FEC approach borrows from the tried-and-true mobile wireless technology: repeat the signal more than once. With mobile wireless systems, this operation interleaves successive copies of the coded voice image across multiple packets (slots in the mobile wireless terminology). With Internet telephony, experiments are underway to send copies of the packet 1 to n times. If one copy is lost, say copy n, it can be recovered from the other copies. However, this operation is only effective if a copy arrives safely, and therefore implies that one of the copies survives the burst error. If the copies are spaced-out too far in time to survive the error, they may arrive too late to be useful.

Order of Arrival of Packets

The subject of the order of the arrival of packets at the receiver is not of keen interest to the data application if it is supported by TCP, because TCP can re-order the TCP segments and present the traffic to the application in the correct order. TCP is not used for voice and video, so

[1][BORE97]. Borella, M. S. "Analysis of End-to-End Internet Packet Loss: Dependence and Asymmetry," *IEEE Network,* Reprint, 1997).

Petitioner's Exhibit 1021
Cisco Systems, Inc. v. Video Solutions Pte. Ltd
Page 000014

the order of packet arrival is an important subject to these applications. As of this writing, studies are underway to capture statistics and discover the incidences of misordered packet arrival. Several studies show that out-of-sequence arrival is not unusual.

Another factor is to note that Internet delays follow a diurnal cycle. During the hours of 8:00 A.M. to 6:00 P.M., delays are greater. For example, in the middle of the business day, delays are about 20 ms greater than in the evenings.

Hop Distance

Hop distance is a term used to describe the number of hops between a sender and a receiver. It is a critical aspect in internet telephony because more hops means more delay, and more variable delay. Hop distance must consider round-trip time (RTT), because of the interactive, real-time nature of telephone conversations. Figure 2–3 shows several aspects of hop distance and RTT, and its deleterious effect on the quality of voice applications.

The traffic is to be sent from a host on subnet 1 to a host on subnet 2. The IP datagrams must be processed by both hosts as well as all the routers on the path between the hosts. Let us assume the traffic traverses through the fewest number of hops (a common approach), which means the datagrams are processed by seven routers, numbered router 1 through router 7 in the figure. Thus, the datagrams are sent through nine hops. If the routers are not heavily loaded with traffic, then queuing delay will be short, and the delay at each router, while variable, will not create a major problem when the traffic arrives at the receiver. However, if traffic is heavy and/or if the routers are not performing their datagram forwarding operations efficiently, the accumulated and variable delay will result in the inability of the receiver to reconstitute the real-time voice signal into a high quality speech pattern.

We learned earlier that round-trip delay in the Internet rarely exceeds 200 ms. But we also know that the delay is highly variable. For example, a delay going through the same number of nodes might be 100 ms on one occasion; on another it may be 200 ms.

However, these studies focus only on the RTT in the network. Keep in mind that this RTT figure does not include analog-to-digital conversion, codec operations, or other factors that would increase the RTT.

Several studies also reveal that it is clear that geographical distance cannot be correlated to round-trip delay. Indeed in one study, a short distance of only 477 miles, but with a hop count of 21 resulted in a 500-ms

Petitioner's Exhibit 1021
Cisco Systems, Inc. v. Video Solutions Pte. Ltd
Page 000015

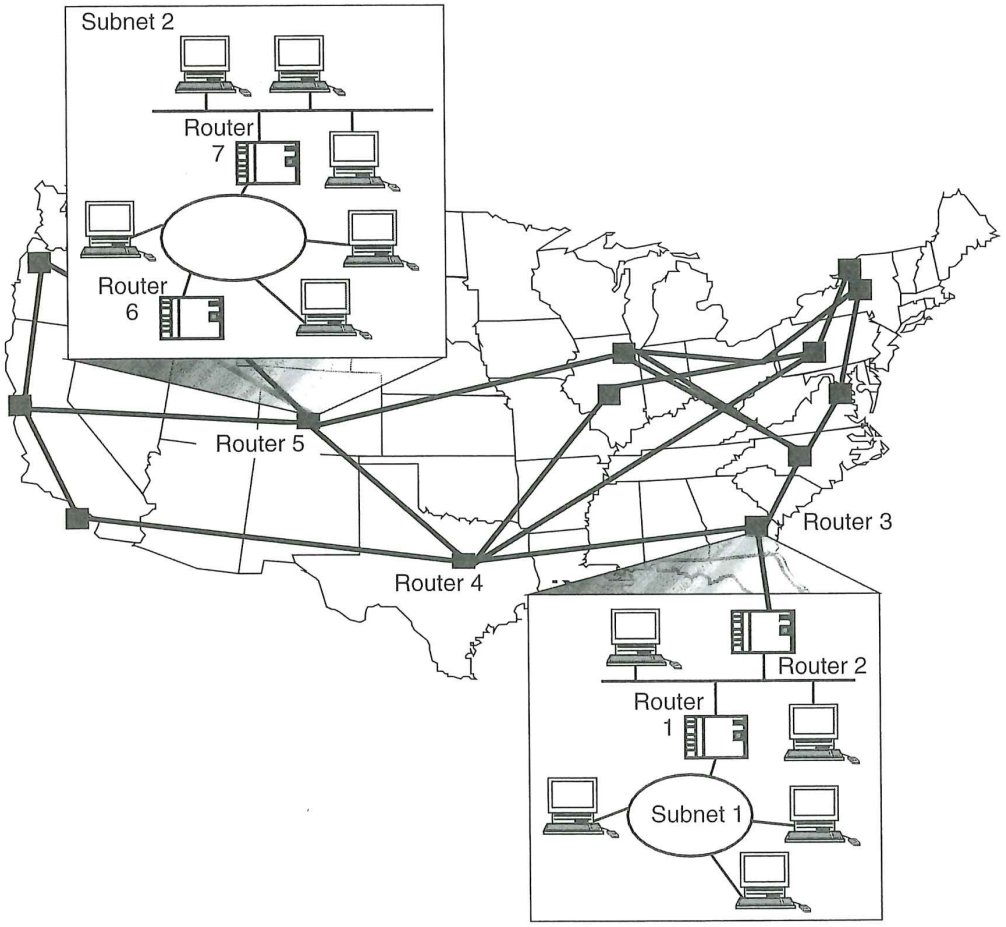


Figure 2–3 Hop distance and round-trip delay (RTT)

round-trip delay. Therefore to emphasize, hop distance is a key factor in delay and geographical distance is less a factor.

The telephone network does not have these problems. First, the path between the talker and listener is fixed during the call set-up. Second, the circuit switches do not queue the traffic. Rather, the voice channels are time division multiplexed into DS0 channels (TDM slots) and sent directly from the input interface on the switch to a corresponding DS0 slot on a preconfigured output interface. The delay through the voice switch is miniscule (it is not even a factor), and fixed.

Thus, circuit switches provide fixed paths, as well as very low, fixed delay. In contrast, packet switches, such as routers, provide variable paths, and variable delay, which is sometimes a low delay and sometimes a lengthy delay.

Petitioner's Exhibit 1021
Cisco Systems, Inc. v. Video Solutions Pte. Ltd
Page 000016

ported with the TCP header. As explained later in this chapter, TCP provides extensive support for traffic integrity operations: (a) error checks, (b) acknowledgments of traffic, (c) retransmissions of errored traffic, and (d) flow control.

But voice applications need not use TCP; in fact, they should not use it, because the TCP support features introduce too much delay in the overall procedure, and it is not practicable to achieve real-time performance if some of the errored packets are retransmitted. The overall delay, would be over 400–500 ms, clearly beyond the RTT that is acceptable for speech.

The potential problem is related to TCP flow-control mechanisms. One feature implemented in many TCP products is the ability of TCP to back off from sending traffic when the Internet is congested and experiencing poor RTT.

In this regard, TCP is a "courteous" protocol. It waits if necessary and does not intrude its presence (its packets) into a congested network. But UDP, on which VoIP operates, has no such qualms. It is a "discourteous" protocol in that is has no back-off mechanisms. It will continue sending traffic even if the network is congested.

So, if the network is congested, TCP stands by the door at the network, and does not enter. Indeed, TCP opens the door, and UDP goes through it!.

As seen in Figure 2–4(b), UDP accounts for only 3 percent of traffic that is placed directly in the IP datagram. However, if voice (and video) become heavy users of the Internet, then it is possible that the TCP users will experience degraded performance during periods of high activity of the real-time traffic.

Earlier, I said it is a *potential* problem. If the Internet and internets are built with sufficient bandwidth, there will be no problem.

OVERVIEW OF IP

Scores of books are available that describe IP in detail. The approach here is to provide an overview of IP, and emphasize the aspects of IP that are especially important for the support of voice traffic. I recommend any of the Douglas Comer and Richard Stevens books to the newcomer to TCP/IP.

IP is an example of a connectionless service. It permits the exchange of traffic between two host computers without any prior call setup. (However, these two computers can share a common connection-oriented trans-

Petitioner's Exhibit 1021
Cisco Systems, Inc. v. Video Solutions Pte. Ltd
Page 000017

port protocol.) Since IP is connectionless, it is possible that the datagrams could be lost between the two end-user's stations. For example, the IP gateway enforces a maximum queue length size; if this queue length is violated, the buffers will overflow. In this situation, the additional datagrams are discarded in the network. For this reason, a higher level transport layer protocol (such as TCP) is essential to recover from these problems.

IP hides the underlying subnetwork from the end user. In this context, it creates a virtual network to that end user. This aspect of IP is quite attractive because it allows different types of networks to attach to an IP node. As a result, IP is reasonably simple to install and because of its connectionless design, it is quite robust.

Since IP is an unreliable, best-effort datagram-type protocol, it has no retransmission mechanisms. It provides no error recovery for the underlying subnetworks. It has no flow-control mechanisms. The user data (datagrams) may be lost, duplicated, or even arrive out of order. It is not the job of IP to deal with most of these problems. As we shall see later, most of the problems are passed to the next higher layer, TCP.

These low-level characteristics of IP translate into a fairly effective means of supporting real-time voice traffic. Assuming the routers are fast, and sufficient bandwidth is available, IP does not introduce significant overhead to the support of VoIP. There are better mechanisms, but as stated in Chapter 1, no other mechanism has the universal presence of IP (and the IP addresses).

Figure 2–5 shows how IP processes an incoming IP datagram [STEV98].[3] The incoming packet is stored in a queue to await processing. Once processing begins, the options field is processed to determine if any options are in the header (the support for this operation varies). The datagram header is checked for any modifications that may have occurred during its journey to this IP node (with a checksum field discussed later). Next, it is determined if the IP address is local; if so, an IP protocol ID field in the header is used to pass the bits in the data field to the next module, such as TCP, UDP, ICMP, etc.

An IP node can be configured to forward or not forward datagrams. If the node is a forwarding node, the IP destination address in the IP datagram header is matched against a routing table to calculate the next node (next hop) that is to receive the datagram. If a match in the table to

[3][STEV98]. Stevens, W. Richard. *TCP/IP Illustrated*, Addison-Wesley, 1997. Mr. Steven's figure in his book (page 112) does not contain the error-check operation, which I have added in this figure.

Petitioner's Exhibit 1021
Cisco Systems, Inc. v. Video Solutions Pte. Ltd
Page 000018

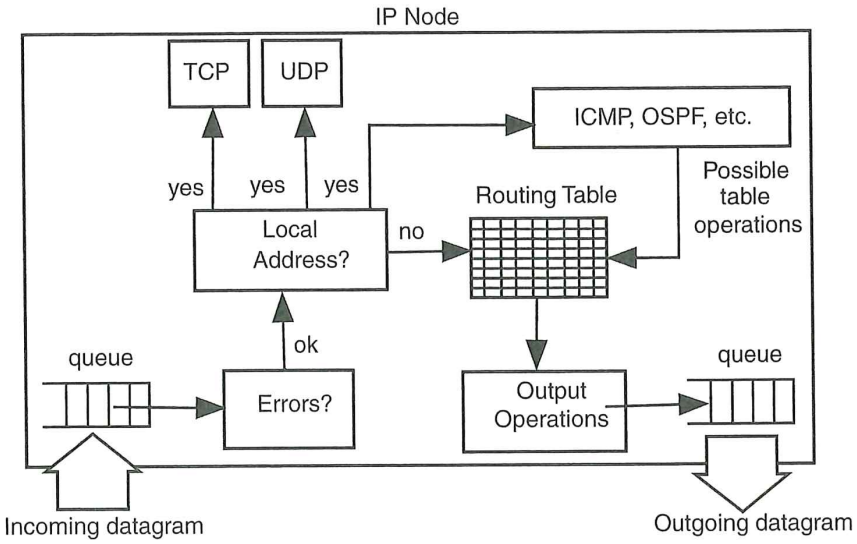


Figure 2–5 Processing an IP datagram

the destination address is found, the datagram is forwarded to the next node. Otherwise, it is sent to a default route, or it is discarded.

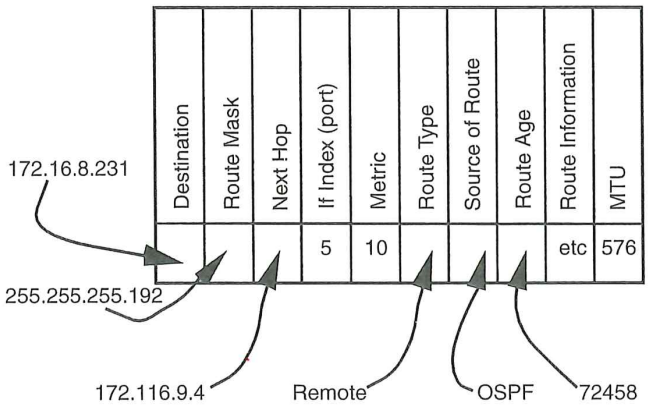
Figure 2–6 is an example of a typical routing table found in a router. Individual systems differ in the contents of the routing table, but they all resemble this example. The entries in the table are:

- **Destination:** IP address of the destination node.
- **Route Mask:** Mask that is used with destination address to identify bits that are used in routing.
- **Next Hop:** IP address of the next hop in the route.
- **If Index (port):** Physical port on the router to reach the next hop address.
- **Metric:** "Cost" to reach the destination address.
- **Route Type:** Directly attached to router (direct), or reached through another router (remote).
- **Source of Route:** How the route was discovered.
- **Route Age:** In seconds, since the route was last updated.
- **Route Information:** Miscellaneous information.
- **MTU:** Size of payload

The IP Datagram

A productive approach to analyzing IP is to first examine the fields in the IP datagram (PDU) depicted in Figure 2–7.

Petitioner's Exhibit 1021
Cisco Systems, Inc. v. Video Solutions Pte. Ltd
Page 000019



Where:
MTU Maximum transmission unit size (in bytes, of the L_2 I field)

Figure 2–6 Typical IP routing table

The *version* field identifies the version of IP in use. Most protocols contain this field because some network nodes may not have the latest release available of the protocol. The current version of IP is 4.

The *header length* field contains 4 bits which are set to a value to indicate the length of the datagram header. The length is measured in 32-bit words. Typically, a header without QOS options contains 20 octets. Therefore, the value in the length field is usually 5.

The *type of service (TOS)* field can be used to identify several QOS functions provided for an Internet application. It is quite similar to the

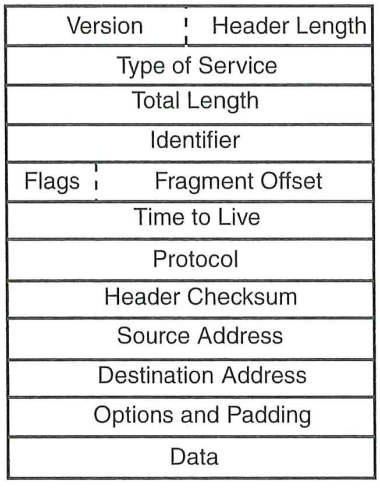


Figure 2–7 The IP datagram

Petitioner's Exhibit 1021
Cisco Systems, Inc. v. Video Solutions Pte. Ltd
Page 000020

Figure 2–8 reinforces those thoughts once more. It shows how the traffic passes from TCP or UDP to the respective application.

A TCP upper-layer user in a host machine is identified by a *port* identifier. The port number is concatenated with the IP internet address to form a *socket*. This address must be unique throughout the internet and a pair of sockets uniquely identifies each endpoint connection. As examples:

Sending socket = Source IP address + source port number
Receiving socket = Destination IP address + destination port number

Although the mapping of ports to higher-layer processes can be handled as an internal matter in a host, the Internet publishes numbers for frequently used higher-level processes.

Even though TCP establishes numbers for frequently used ports, the numbers and values above 1024 are available for private use. The remainder of the values for the assigned port numbers have the low-order 8-bits set to zero. The remainder of these bits are available to any organization to use as they choose.

Examination of port numbers. Recently, routers have begun examining the port numbers in order to determine the type of traffic in the user payload. For example, a destination port number might identify a video application and the router could then treat this traffic as a high-priority unit (in contrast to low-priority data traffic).

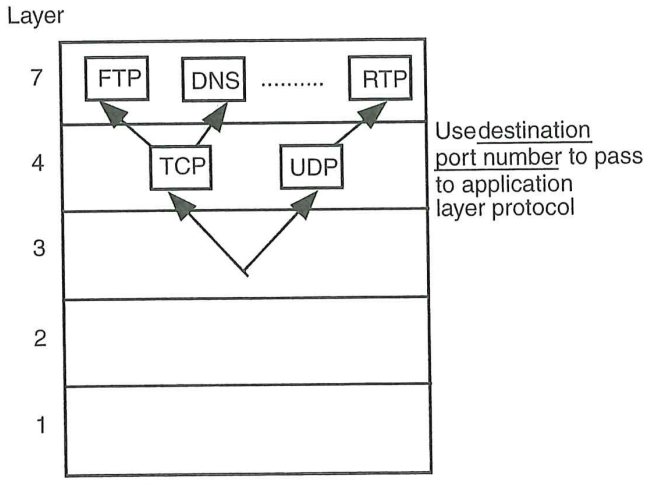


Figure 2–8 The port concept

TCP Traffic Management Operations

TCP acknowledges segments for both interactive traffic and bulk data traffic. However, TCP varies in how these acknowledgements occur. The variability depends on the operating system through which TCP executes, as well as the options chosen for a TCP session. In essence, there is no "standard" TCP operating profile.

Notwithstanding these comments, TCP does use many common procedures that are found in most TCP implementations and versions.

UDP

The User Datagram Protocol (UDP) is classified as a connectionless protocol. It is sometimes used in place of TCP in situations where the full services of TCP are not needed. For example, the Trivial File Transfer Protocol (TFTP) and the Remote Procedure Call (RPC) use UDP.

UDP serves as a simple application interface to the IP. Since it has no reliability, flow control, nor error-recovery measures, it serves principally as a multiplexer/demultiplexer for the receiving and sending of IP traffic.

UDP uses the port concept to direct the datagrams to the proper upper-layer application. The UDP datagram contains a destination port number and a source port number. The destination number is used by UDP and the operating system to deliver the traffic to the proper recipient.

Perhaps the best way to explain UDP is to examine the message and the fields that reside in the message. As Figure 2–9 illustrates, the format is quite simple and contains the following fields:

- *Source Port*: This value identifies the port of the sending application process. The field is optional. If it is not used, a value of 0 is inserted in this field.
- *Destination Port*: This value identifies the receiving process on the destination host machine.
- *Length*: This value indicates the length of the user datagram including the header and the data. This value implies that the minimum length is 8 octets.
- *Checksum*: This value is the 16-bit one's complement of the one's complement sum of the pseudo-IP header, the UDP header, and the data. It also performs a checksum on any padding (if it was necessary to make the message contain a multiple of two octets).

Petitioner's Exhibit 1021
Cisco Systems, Inc. v. Video Solutions Pte. Ltd
Page 000022

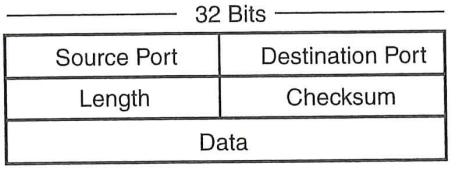


Figure 2–9 Format for the UDP data unit

There is not a lot more to be said about UDP. It is a minimal level of service used in many transaction-based application systems. However, it is quite useful if the full services of TCP are not needed.

SUMMARY

Data networks in general, and the Internet in particular, are designed to support asynchronous, non–real-time traffic. The variable delay in data networks is usually not a big problem, due to the asynchronous nature of data applications.

The design philosophy of data networks is the opposite of voice networks. Voice networks are designed to support synchronous real-time traffic, with a fixed delay between the sender and the receiver.

Three obstacles must be overcome if the Internet and internets are to be effective bearers of speech: (a) latency must be reduced, (b) jitter must be reduced, (b) traffic loss must be improved.

UDP and IP are the chosen L_4 and L_3 protocols for VoIP. Their choice has nothing to do with their merit (although they are not poor choices). Rather, it reflects the fact that UDP and IP are found in most workstations and personal computers.

Petitioner's Exhibit 1021
Cisco Systems, Inc. v. Video Solutions Pte. Ltd
Page 000023

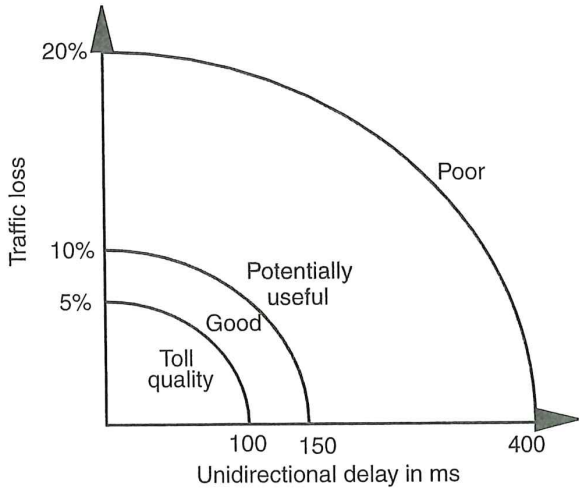


Figure 8–1 Traffic loss versus traffic delay [COX98]

quality. One could take this idea to the extreme and postulate a packet of say 1 byte (or even 1 bit!). Obviously, there is a point of diminishing returns where the overhead of headers to the miniscule user packet is so high that it militates against building an efficient network.

Another consideration is to attempt to build the packet size in consonance with the output of the codec. For example, the G.729 and G.729a codecs produce a 10-ms sample (10 bytes, with 1 byte per ms). In so far as possible, it is a good idea to package the packets in 10-ms bundles.

But this approach may not be possible. For example, if these voice packets are placed inside a fixed-length frame, they may not fit exactly in the fixed payload of the frame. If this is the situation, and the payload size of the frame cannot be altered, the system should be able to break the sample into even-byte boundaries and place it into successive frames in order to best use the payload bandwidth. This approach is common today, and I provide some examples in Chapter 12.

PERFORMANCE OF VoIP IN PRIVATE SYSTEMS

Our focus for this part of the chapter is to examine two studies conducted by Mier Communications Inc. on VoIP products. The general results of the first study are available from Mier Communications

Petitioner's Exhibit 1021
Cisco Systems, Inc. v. Video Solutions Pte. Ltd
Page 000024

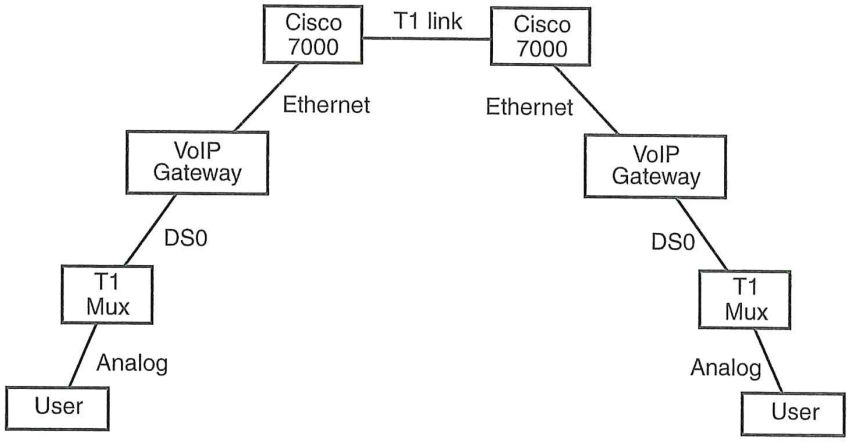


Figure 8–2 Topology for study [MIER98]

[MIER98].[2] The second study is also available from [MIER99] and [MIER99a]. Mier Communications can be reached at ed@mier.com. After we examine the first test, we then look at the new test, and compare the differences.

Figure 8–2 shows the topology layout and configuration for the first study conducted by Mier Communications. The voice traffic was exchanged between two Cisco 7000 routers over an unchannelized T1 link. The voice traffic was created through a simulated analog line and fed into a T1 mux where it was digitized into 64 kbit/s DS0 slots. This signal was fed into the voice-over-IP (VoIP) gateway (the products under test) where they were then input into the Cisco router across an Ethernet 10 Mbit/s channel. The signals were then transported across the unchannelized T1 link to the receiving router where the signal then was sent through an Ethernet interface to the receiving VoIP gateway, then via a DS0 slot to a T1 mux. The mux converted the signal back to analog and transported the signal to the user at either a telephone handset or a speaker.

The study was performed with various voices (female speakers, male speakers, etc.). All speakers recited common voice images such as *do*, *re*, *me*, and so on.

---

[2][MIER98] Mier, Edwin, E. "Voice-Over-IP: Sounding Better," *Business Communications Review*, February, 1998.

[MIER99a] Mier, Edwin, E. "VoIP Gateways-Tradeoffs Affecting Voice Quality," *Business Communications Review Voice 2000*, January, 1999.

Petitioner's Exhibit 1021
Cisco Systems, Inc. v. Video Solutions Pte. Ltd
Page 000025

Keep in mind that this study did not test voice quality over the Internet. The test was run over the T1 link between the routers. Consequently, you should not infer that these tests state anything about the performance of the Internet in supporting telephony. But the study does provide an interesting and useful assessment of VoIP products. As we shall see, the test reveals that high-quality products are available to run voice over what is essentially data-based protocols, and offer an attractive alternative to public-switched toll services.

Several tests were performed on four products. According to Mier Communications, all four products performed well and provided high-quality telephony images to the end-user. The four products tested were Lucent Technologies, Micom (now part of Nortel), Nuera Communications, and Selsius Systems.

Figure 8–3 reflects the assessment of the quality of these systems when operating under ideal network conditions, The definition of ideal network conditions is a fully available T1 channel which is then unchannelized (not restricted to DS0, TDM slots). In addition, the media operated under relatively error-free conditions without adding delay or latency across the system. That is to say, any latency or delay in the test

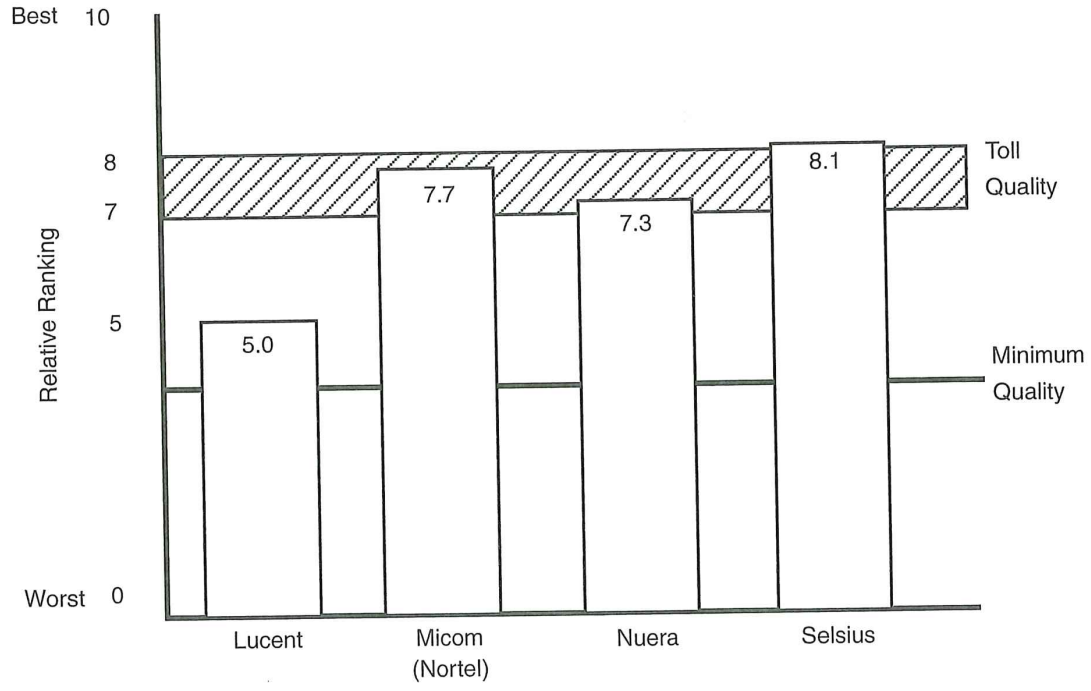


Figure 8–3 Voice quality under ideal conditions [MIER98]

Petitioner's Exhibit 1021
Cisco Systems, Inc. v. Video Solutions Pte. Ltd
Page 000026

was introduced as a result of the components in the topology and the operations of the VoIPs.

Obviously, the evaluation of the tests is subjective, but the evaluation of voice quality on the telephone network is subjective as well. Notwithstanding, as this figure shows, all products were evaluated as exceeding a minimum quality for the signal. The scale of 0 to 10 was devised with the expectation that a score between 7 and 8 represented the current "toll quality" exhibited in the telephone network. The evaluation of 10 would reflect the best possible telephone voice connection wherein the analog loop is terminated into a digital system one time only, sent a very short distance, and then converted back to an analog signal.

Next, the test was conducted in an environment that exhibited poor quality conditions. To simulate these poor conditions, Mier Communications applied burst errors to the T1 line. The approach was to cycle a 5-ms burst of random errors at a rate of 1×10^2, followed by 50 ms of transmissions with a bit error rate of 1×10^6. The 1×10^6 value is a reasonable assumption for BER performance on a typical local loop. The result of this operation is the introduction of errors in about 10 percent of the IP telephony datagrams. As Figure 8–4 shows, Micom and Nuera

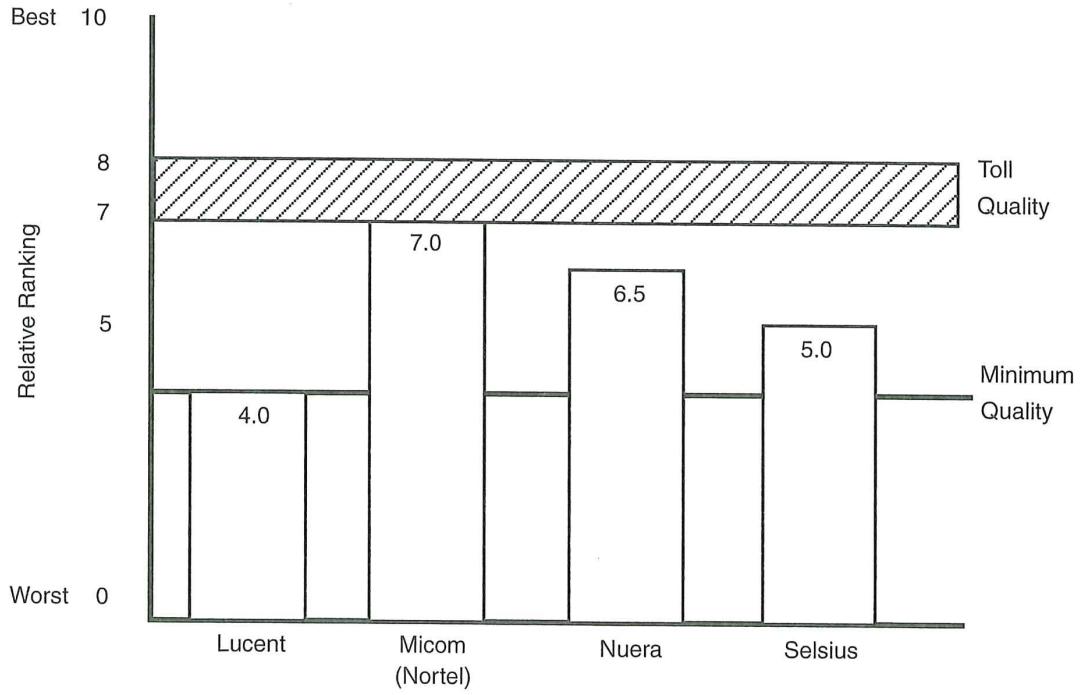


Figure 8–4 Voice quality under poor conditions [MIER98]

Petitioner's Exhibit 1021
Cisco Systems, Inc. v. Video Solutions Pte. Ltd
Page 000027

continue to exhibit high-quality performance and Lucent and Selsius still exhibited somewhat acceptable quality.

Figure 8–5 shows the bandwidth utilization of the four tested products working under relatively error-free conditions and on a full-unchannelized T1 link. The Nuera system scored the highest with a bandwidth requirement of 14 kbit/s followed by Micom with 18 kbit/s. The Nuera performance approaches that of the sophisticated mobile/wireless vocoders that operate in a range of 13 kbit/s.

Not shown in this figure is the underlying fact that voice quality improved with those systems that used smaller protocol data units (with one exception). Micom's packet size is 91 bytes; Nuera's packet size is 93 bytes; Lucent's packet size is 77 bytes; and Selsius system uses 300 bytes.

In fairness, it should be stated that the performance of Nuera and Micom is partially attributable to the fact that they are using proprietary vocoding techniques whereas Lucent employs the ITU-T G.723.1 and Selsius uses ITU-T G.711. All organizations that belong to the VoIP and the International Media Teleconferencing Consortium (IMTC) have all selected G.723.1 for their basic vocoder.

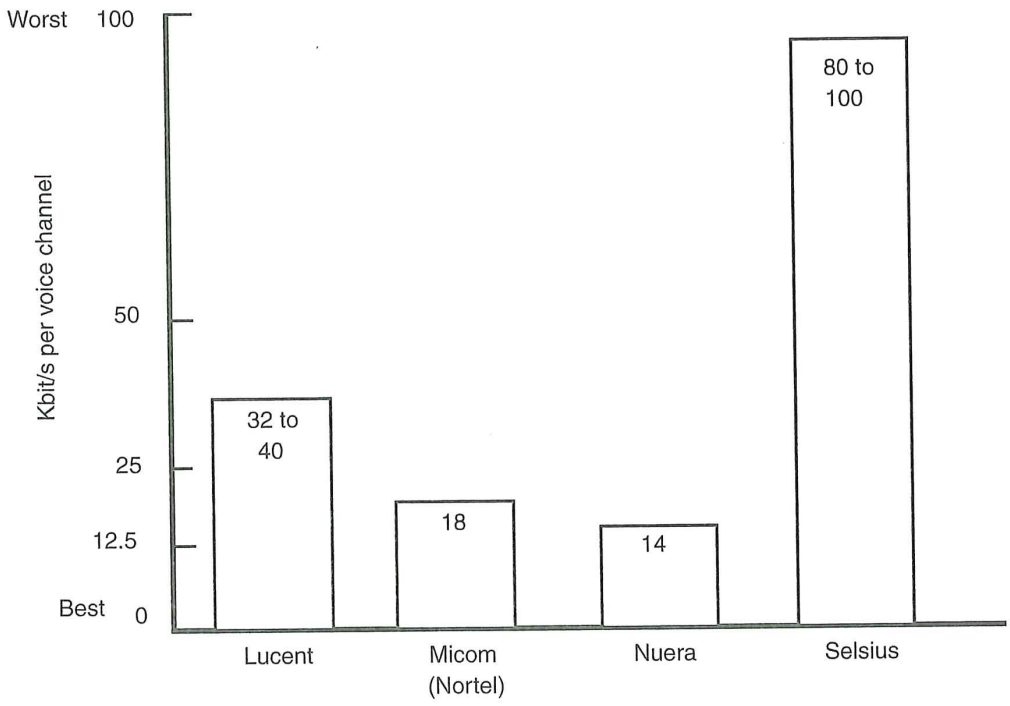


Figure 8–5 Bandwidth used under ideal conditions [MIER98]

Petitioner's Exhibit 1021
Cisco Systems, Inc. v. Video Solutions Pte. Ltd
Page 000028

It is reasonable to assume that the Lucent and Selsius performance would improve with proprietary schemes.

The trade-off of proprietary versus standardized schemes pertain to the fact that private systems might opt for proprietary schemes because of their superior performance. But if different vendor systems are to interoperate with each other, the standards must be used.

Figure 8–6 shows the latency introduced by the tested systems. All operated within the 250-ms threshold (which is not considered very good performance), and none operated in accordance with the telco standard for one-way latency of 100 ms. The 250-ms delay is considered by some to be the threshold at which it becomes noticeable and disturbing to the users of the system.

In addition, the general study did not focus on where the delay was encountered. Recent tests with voice-over-telephony, in some products (not those cited in the study), indicate that significant delay is occurring in the line card located at the customer premises. Therefore, a clear analysis must focus on the specific components that create the delay because the VoIP may not be the actual culprit for these delay performances.

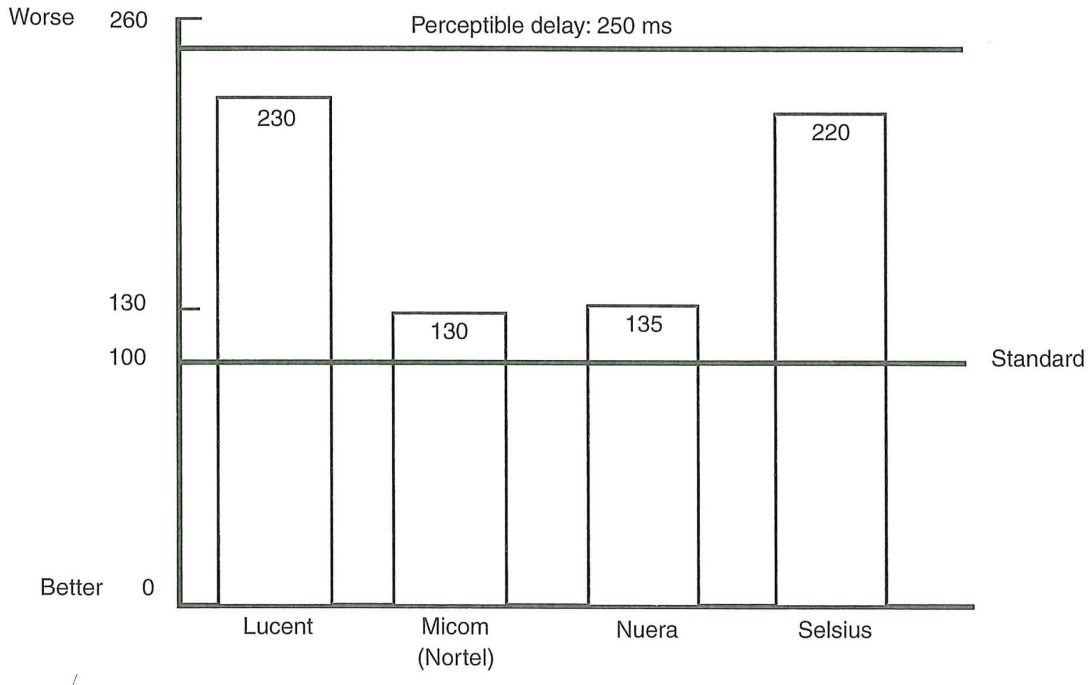


Figure 8–6 Latency: One-way delay [MIER98]

Petitioner's Exhibit 1021
Cisco Systems, Inc. v. Video Solutions Pte. Ltd
Page 000029

As a result of the interest shown in Mier's first test, the company conducted another test [MIER99]. You can reach Mier at (609) 275–7311 if you need more information about these studies.

Figure 8–7 shows the layout for the study, and notations of the vendors' testing equipment (the tested systems are not shown here). Several points should be emphasized here. First, the test layout is more elaborate and sophisticated than the earlier test. Second, G.723.1 and G.711 were not used by the vendors in this test. Instead, G.729 and G.729a were used. Third, H.323 was part of a test, and it is considered an important component of VoIP by all the tested vendors. Fourth, the results of the test were considerably higher than the earlier test, reflecting better coders and the overall maturation of the technology.

The analog voice files were prerecorded and sent through an Adtran TSU/100 T1 mux, where they were converted to DS0 signals, then to a

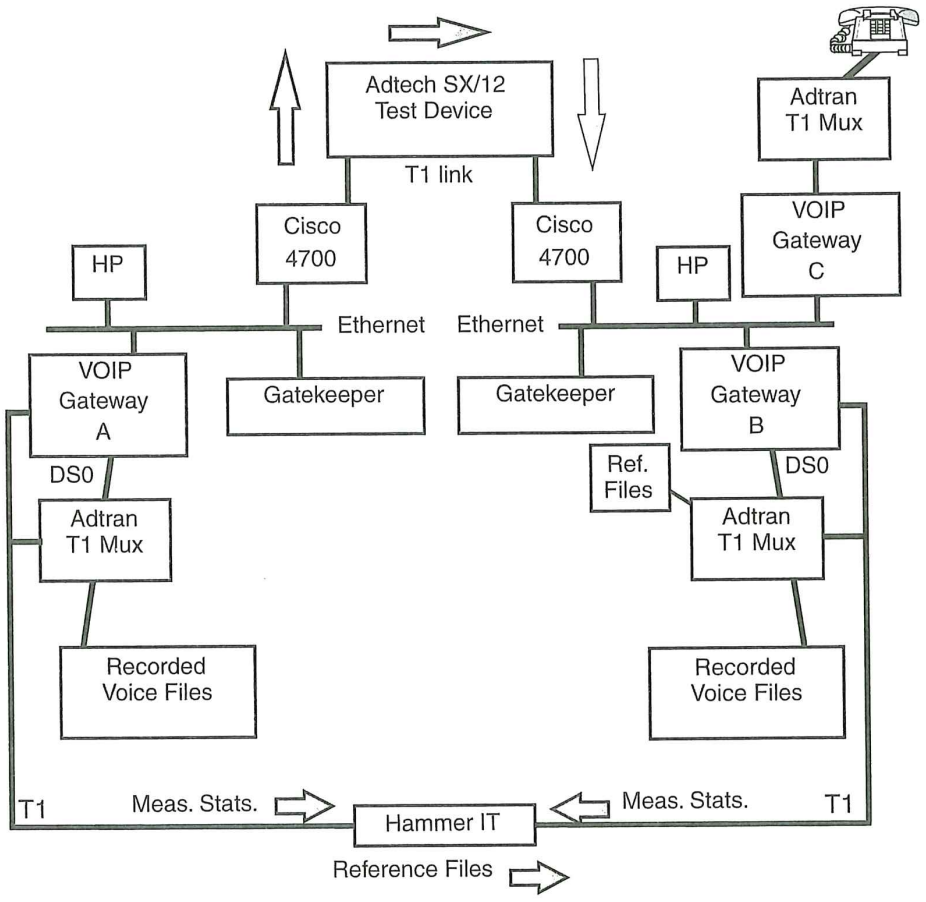


Figure 8–7 Revised test topology [MIER99]

VoIP gateway (Gateway A). Here they were subjected to the vendor's codec and VoIP operations. The resultant VoIP packets were sent through a Cisco 4700 router across an unchannelized T1 (and through an Adtech SX/12 test device) to the receiving site. At this site, the process was reversed (Gateway B), with the voice message rerecorded onto the voice files. A back-to-back T1 line was used to provide reference signals at the receiver. These toll quality samples were compared (by a panel of people) against the tested samples. Gateway C was used to test real-time quality and latency between two people.

The original voice files were sent to the receiver to act as reference (ref.) files for the comparison tests. A scope from Fluke Corp. was used to measure the latency that each vendor's product added (not shown in the figure). A Hammer IT VoIP Test System also measured latency and generated loads on the gateways, as well as measured voice quality and a variety of other statistical support.

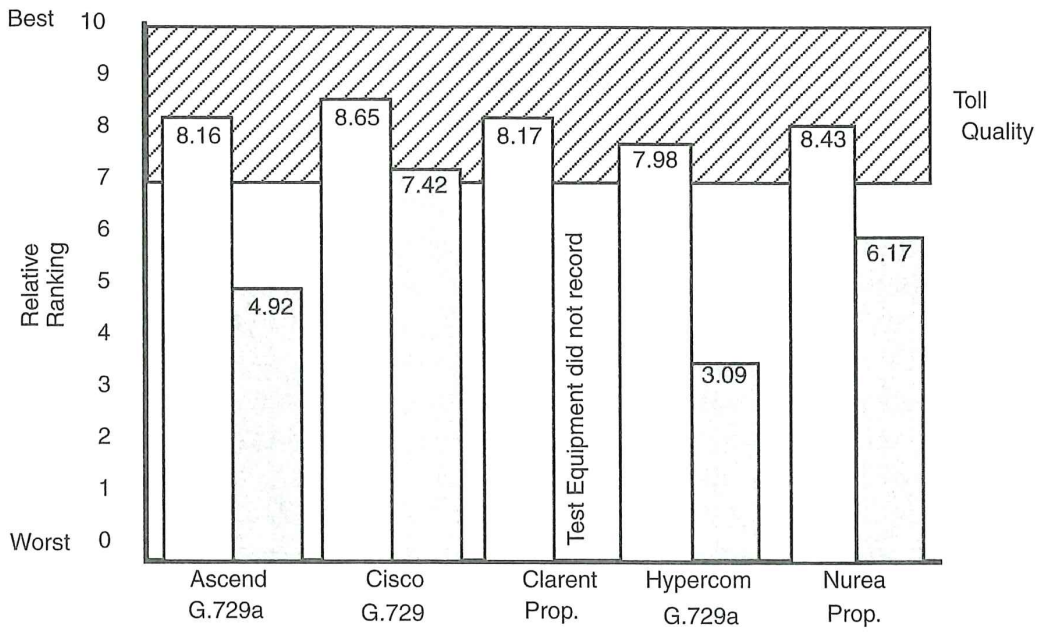
The testing for voice quality was performed by Mier's staff by listening to the rerecorded voice samples, as well as using interactive tests between two people. As Figure 8–8(a) shows, all vendors' products performed very well, under ideal conditions (the higher bars for each vendor). All scored well above minimum toll quality. Mier's conclusions are that the reception of packet voice with the tested VoIP gateways cannot be distinguished from POTS services. The lower bars in Figure 8–8(a) show the quality during poor conditions, in which periodic error bursts were placed on the link.

The figure also shows the codecs used for the tests. The vendors used G.729, G.729a, or proprietary coders. All vendors have products with other codecs, such as G.711, and most of them have or are planning to support G.723.1.

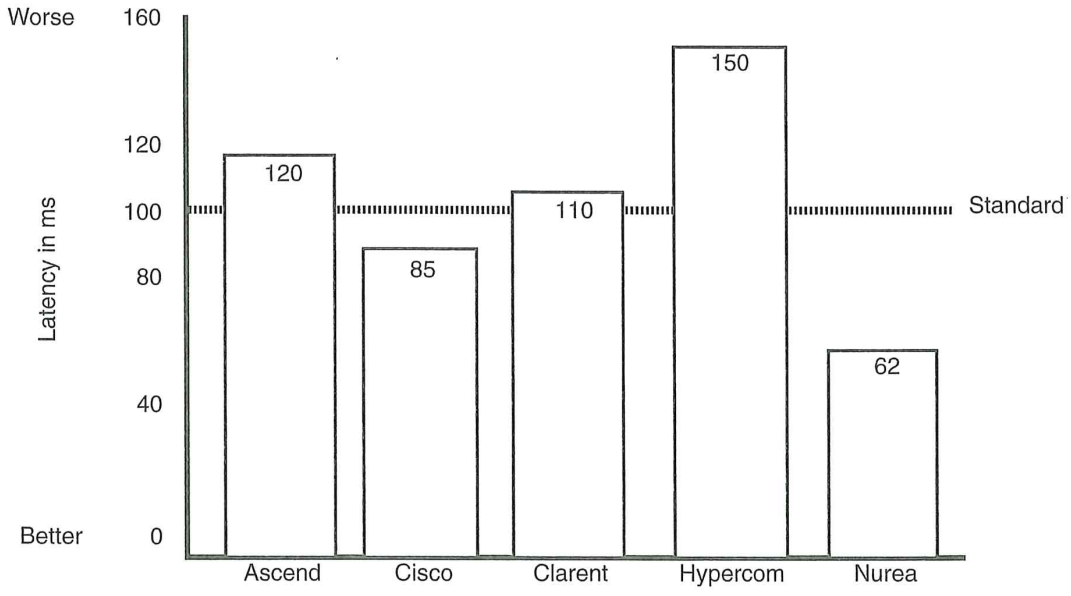
Figure 8–8(b) shows the one-way latency for the tested products. Mier's ongoing tests have made them attuned to quality vs. latency, and they conclude that once latency drops below 90 to 80 ms, a person cannot tell the difference between the vendors, or for that matter, between the vendors and the public switched network. After all the tests, the testers could actually tell the differences between latencies under 100 ms, slightly more than 100 ms, and considerably more than 100 ms.

This study was amplified further in [MIER99a]. As you might expect if you read the first part of this chapter, the vendors whose traffic experienced the lowest latency also experienced the best quality. What you might also be thinking is how large a packet do the vendors' products produce? After all, we learned that larger packets lead to more loss on noisy links, and they increase latency.

Petitioner's Exhibit 1021
Cisco Systems, Inc. v. Video Solutions Pte. Ltd
Page 000031



(a) Voice quality tests



(b) One-way latency tests

Figure 8–8 Voice quality and latency [MIER99]

Then it should come as no surprise that the vendors (with some minor exceptions) whose product gave the best voice quality also exhibited the lowest latency. The number and size of the voice samples in each packet are shown in Figure 8–9.

From these tests, it seems prudent to construct a small packet if one is interested in voice quality. Indeed it is, but there is a price to pay. As stated in the introductory part of the chapter, the packet that contains a small number of samples suffers from more overhead than a packet that contains more samples. This is so, because number of bytes in the associated headers (protocol control information) remain the same for a small or larger user payload—the IP header is always 20 bytes, etc.

Mier's study corroborates these claims, as shown in Figure 8–10. The bandwidth consumption was tested on the LAN and a simulated WAN. Statistics were gathered on the LAN by an HP Internet Advisor, and on the WAN by an Adtech SX-12. In addition, on the WAN, the Adtech gradually reduced the amount of bandwidth available (to the point where the quality was unacceptable). As expected, the systems that

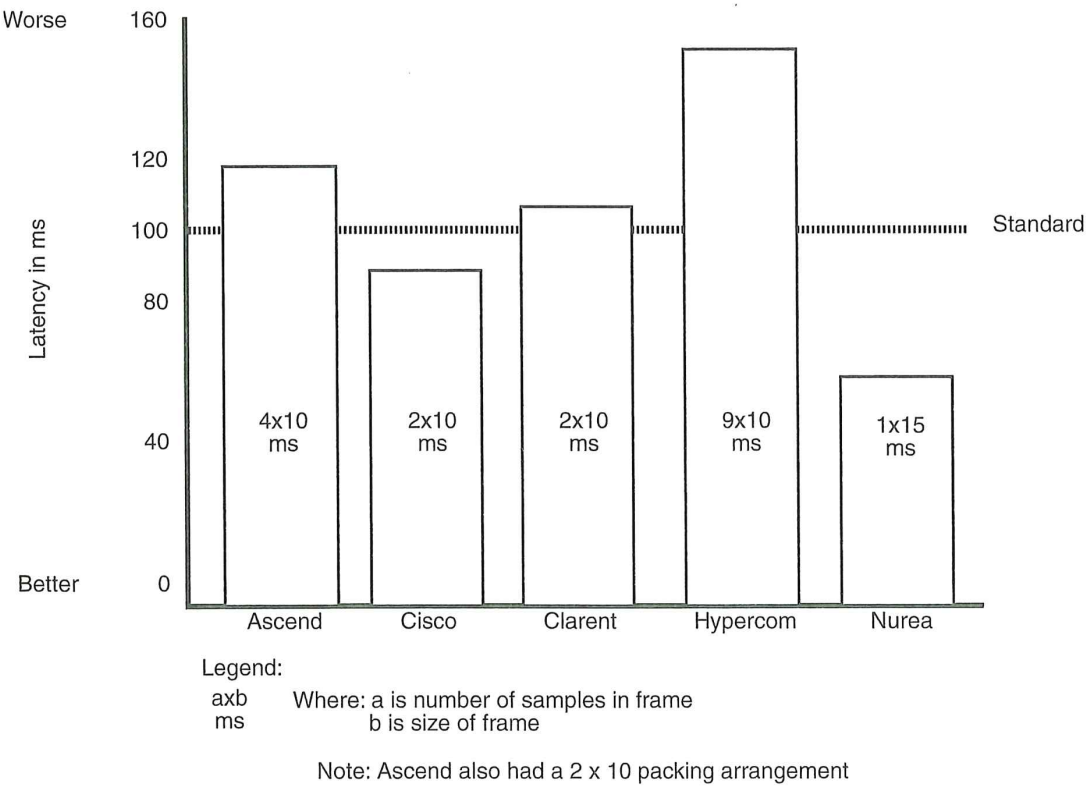


Figure 8–9 Latency and frame size [MIER99]

Petitioner's Exhibit 1021
Cisco Systems, Inc. v. Video Solutions Pte. Ltd
Page 000033

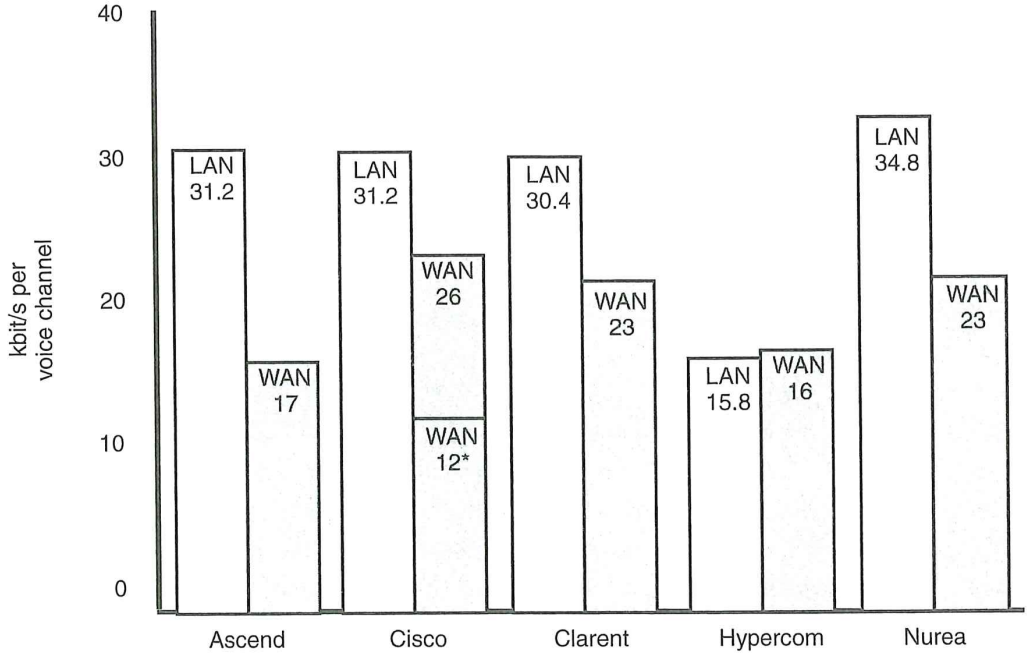


Figure 8–10 Bandwidth consumption [MIER99]

had larger packet sizes and more latency also had a more efficient use of bandwidth.

As stated in the first part of this chapter, there is a tradeoff of latency, packet size, and perceived voice quality. If bandwidth is plentiful, one would opt for the system that provides higher quality, and so on.

In addition, these vendors' gateways are configurable. Frame packing (number of samples per frame) can be set on most of them, as can the jitter buffer. Many other factors come into play in evaluating VoIP gateways, such as costs and ease of use (configuration). So, it is best to delve into more detail about a vendor's product, but the Mier studies are excellent places to start.

PERFORMANCE OF VoIP IN PUBLIC SYSTEMS

This part of the chapter examines a test conducted by 3Com on the performance of VoIP in the public Internet. The source for this study is [COX98].

The study entailed sending and receiving traffic between three nodes: University of California, Davis; University of Illinois, Chicago;

Petitioner's Exhibit 1021
Cisco Systems, Inc. v. Video Solutions Pte. Ltd
Page 000034

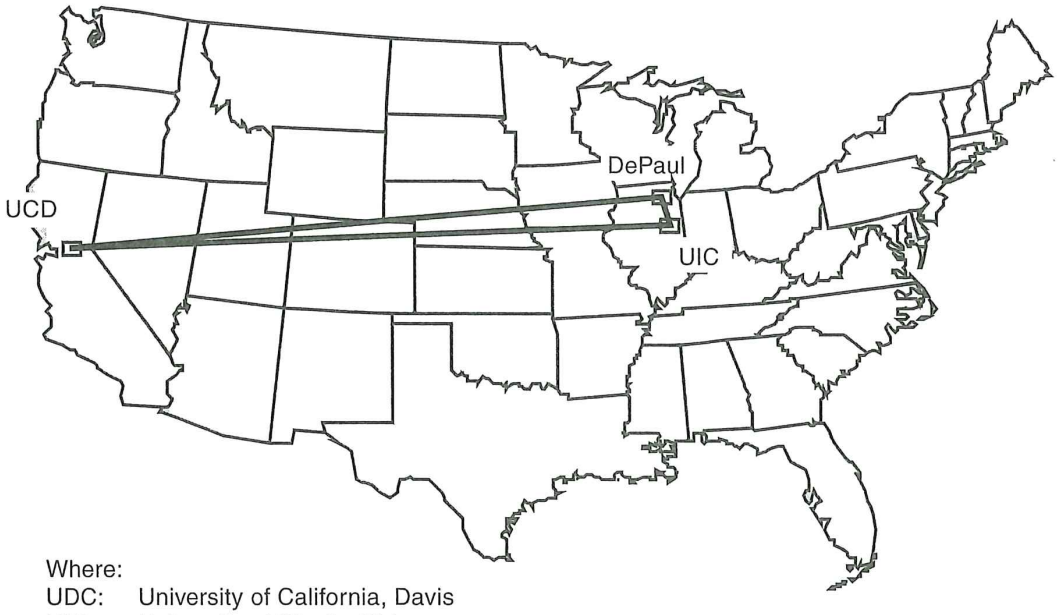


Figure 8–11 Topology for the study [KOST98][3]

and DePaul University, see Figure 8–11. The tests were run for a six-month period. During these tests, a client would transmit once per hour to a server for three minutes. The transmissions involved a trace which allowed the analysts to judge RTT as well as packet loss. The engineers designed their own "ping" program and did not rely on internet pings in order to control how routers handle conventional ping packets. Observations were made in the evenings as well as various times during the business day, and on weekends.

In addition, different codecs were employed in the tests with the emphasis on G.723.1 and G.729A. Tests were done using PC-to-PC communications and VoIP gateways.

We can summarize several key aspects of the study with the examination of the three figures depicted in Figure 8–12. Figure 8–12 (a) compares the average RTT (in ms) in relation to the hop count. The hop count represents the number of nodes traversed between the client and the server. This figure reveals some interesting facts. The first fact is that RTT exceeds 200 ms. This conclusion is borne out by other studies.

[3][KOST98] Kostas, T. J., Borella, M. S., Sidhu, I., Shuster, G. M., Grabiec, J., and Mahler, J. of a 3Com, paper "Real-Time Voice Over Packet Switched Networks, *IEEE Network,* January, February 1998.

Petitioner's Exhibit 1021
Cisco Systems, Inc. v. Video Solutions Pte. Ltd
Page 000035

The second fact is that the delay is highly variable. On occasion a delay going through the same number of nodes is, say, 100 ms and on another occasion it may be 200 ms.

Keep in mind that these data represent RTT measured with a ping only and does not include analog-to-digital conversion, codec operations, or other factors that would increase RTT.

Figure 8–12(b) is the same figure as that of 8–12(a). The oval is placed on the figure to emphasize that despite the variability in the RTT in relation to hop count, increased hop counts do indeed contribute to delay.

Figure 8–12(c) has some numbers placed around several of the points in the graph. These numbers represent the geographical distances in miles between the sender and receiver of the trace. It is clear that geographical distance cannot be correlated to RTT. Indeed, a short distance of only 477 miles with a hop count of 21 resulted in a 24-ms RTT. Therefore to emphasize, hop distance is a key factor in delay, and geographical distance is less a factor.

Figure 8–13 provides a summary of the 3COM study. Based on the two alternatives of (a) use of the telephone and VoIP gateways, and (b) use of PCs and routers, the study reveals that the telephone/gateway approach provides significantly better performance than the PC/router approach.

Under ideal conditions, option (a) meets the RTT established by the ITU-T specifications. Option (b) does not meet these requirements, but comes close. During less-than-ideal conditions, neither approach meets the requirements, but a large segment of the population would likely find the performance acceptable for option (a). Option (b) pushes the envelope of acceptable quality.

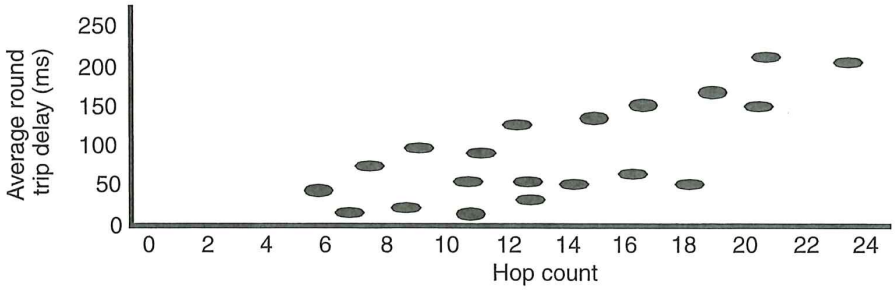
So, is Internet telephony feasible? Yes, but under the present Internet environment, it may not be acceptable to some people, and may be acceptable to others.

Nonetheless, given the attractive features of Internet telephony (one link to the home, integrating voice and data, and low costs), it will surely succeed.

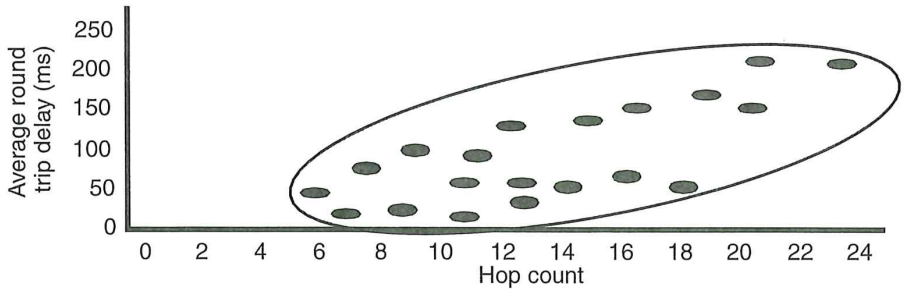
One other aspect of this subject bears examination: the deployment of high-speed (a) ADSL modems, (b) cable modems, or (c) fixed wireless access technologies on the local loop. Once the customer has these technologies available, the equation changes.

First, overhead (headers and trailers) is not as significant a factor, since the increased bandwidth can support this overhead. Second, new PCs will be upgraded to support faster voice coders to take advantage of

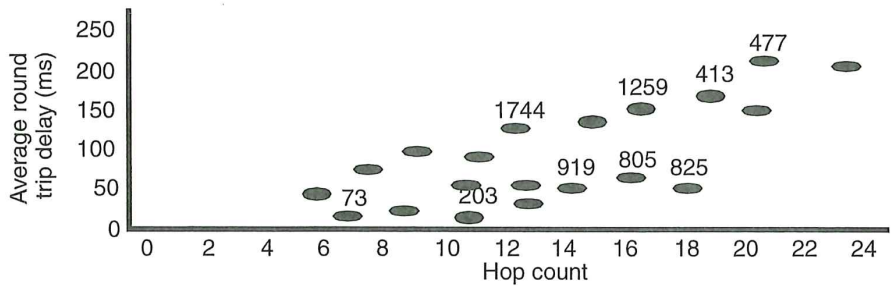
Petitioner's Exhibit 1021
Cisco Systems, Inc. v. Video Solutions Pte. Ltd
Page 000036



(a) Average delay and hop count



(b) Correlation of delay to distance



(c) Relationship of delay and geographical distance

(note: numbers represent distance in miles)

Figure 8–12 Round trip delay versus hop count [KOST98]

the higher-speed local loop. Third, the increased pipes into and out of the Internet will force an upgrading of the Internet's capacity. Fourth, the increase of voice (and video) traffic will also force the Internet to look more and more like the telephone network, but with significantly enhanced multiapplication capabilities.

Petitioner's Exhibit 1021
Cisco Systems, Inc. v. Video Solutions Pte. Ltd
Page 000037

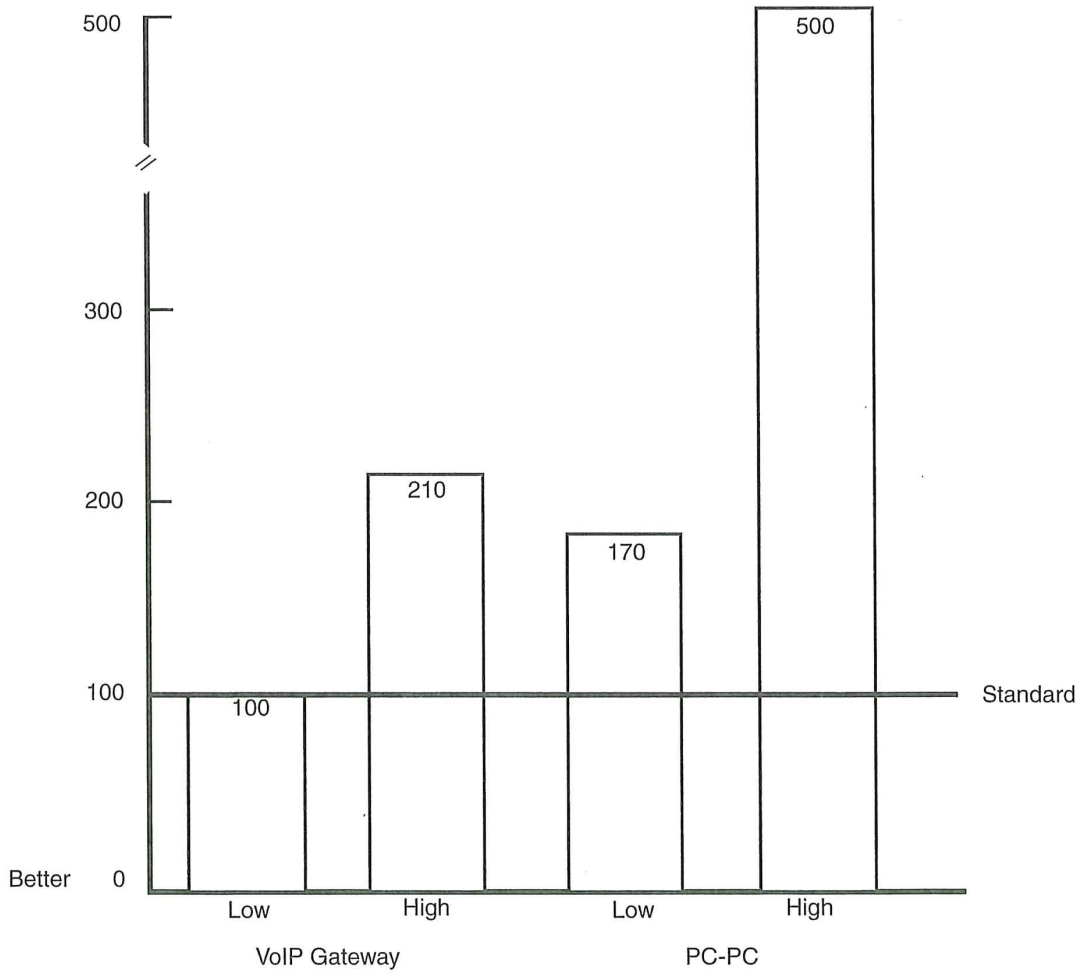


Figure 8–13 Unidirectional delays [KOST98]

SUMMARY

The Mier and Cox studies demonstrate that VoIP is quite feasible in private networks with leased lines, but marginal in the public Internet. It is a good idea to remember that packet voice is in its infancy, and the test results from Mier's labs are impressive. It is reasonable to expect that the Internet will also deliver toll-quality speech—in a few years.

Petitioner's Exhibit 1021
Cisco Systems, Inc. v. Video Solutions Pte. Ltd
Page 000038

Internet Telephony
Voice Over IP

VOICE OVER IP

UYLESS BLACK

The authoritative guide to Internet telephony

Managing and optimizing the performance of IP voice networks

Configuring IP-based call centers

Internetworking IP with SS7

Uyless Black has written the essential guide for telecommunications professionals who must understand voice over IP (VoIP)—or deploy it. Clearheaded and free of hype, *Voice Over IP* carefully evaluates both VoIP's challenges and its compelling advantages, and then reviews each technical standard and critical issue associated with successful deployment.

Start by reviewing the key Internet and IP characteristics that make VoIP difficult, including packet loss and variable delay. Next, understand the role of digital signal processors (DSPs) and voice coders in VoIP. Learn how to establish paths to service providers through the local loop via ISDN, xDSL, HFC, or other approaches; review modem technology for VoIP applications; and understand today's key Internet telephony protocols.

Coverage includes:

- Integrating total VoIP solutions: DSPs, voice coders, high-speed modems, gateways, local loop connectivity, and more
- The role of VoIP gateways and gatekeepers
- Internetworking SS7 with IP and H.323
- VoIP supporting protocols for multicasting (IGMP and MBONE), bandwidth reservation (RSVP, RTP, RTCP), and security services

Black describes today's best approaches for managing performance in both private and public IP networks, compares VoIP with packet voice alternatives such as Voice over Frame Relay (VoFR) and Voice over ATM (VoATM), and more. You'll also find convenient references to telephony signaling, ISDN and SS7, and V.90 modems.

Every telecommunications professional will be impacted by the VoIP revolution. Whether you're evaluating or deploying VoIP, this book places a world-class telecom consultant at your side, delivering all the objective information and insight you need to succeed.

About the Author

Uyless Black is a widely known and respected consultant and lecturer on computer networks and data communications. He is author of all the books in the Prentice Hall Series in Advanced Communications Technologies, including *Residential Broadband Networks, ATM, Vol. I: Foundation for Broadband Networks,* and *Advanced Features of the Internet.*

PRENTICE HALL
Upper Saddle River, NJ 07458

http://www.phptr.com

ISBN 0-13-022463-4

90000

9 780130 224637

Petitioner's Exhibit 1021
Cisco Systems, Inc. v. Video Solutions Pte. Ltd
Page 000039