Quality Management: Troubleshooting Techniques for Voice over IP

Overview

This paper is intended to offer an overview of Voice over IP (VoIP). It will discuss the technical parameters related to ensuring good quality voice output. It will also provide information about measuring the critical parameters that affect VoIP and techniques for measuring these parameters using Fluke Networks' Protocol Expert and Link Analyzer. There is also a section that will briefly describe some "best practice" troubleshooting techniques, using Protocol Expert and Link Analyzer. Industry standards for the measurement of voice quality will be described, along with examples of how network problems degrade the quality of voice.

Packetized voice

VoIP makes it possible to represent an audio/voice signal within a digital bit stream that can be placed in IP packets and carried over the Internet or a corporate intranet. Typical implementations of VoIP are based on collecting 10 ms (1/100 second) of voice, consisting of 80 samples, each representing 1 byte. Usually two such blocks are placed in a data packet for transmission. In order to have a standards-based implementation, a protocol structure is needed. The IP, UDP and real time (RTP) protocols deliver these voice packets between sender and destination, as described in Figure 1.

In order to understand the characteristics of VoIP, it is important to consider the packet networks that support packetized telephony.

Each packet is transmitted through a packet network in a repetitive process. The sending station builds the packet, gains access to the network and sends the packet to the first relaying station, such as a switch, router or gateway. When that device receives the packet, it may run an error

Figure 1

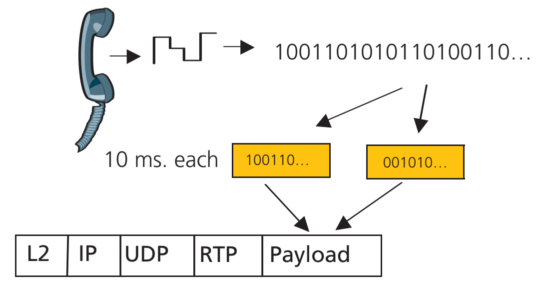

check to see if it is corrupted. If the packet is corrupted, the device will generally drop it. If the packet is not corrupted, the device sends the packet to the next relaying device, repeating the process. No two packets may experience the same delays, nor will they necessarily follow the same route. Packets may also be queued or dropped due to a congested network. As a result, packets representing the same original voice message may arrive erratically, at different times, out of order, or fail to arrive at all.

By comparison, in conventional circuit switched networks, voice bits are transmitted in order, normally at 64kb/sec. They experience comparatively few delays and always follow the same path. That is why the speaker (voice message sender) and listener (voice message receiver) often believe they are connected by a pair of wires on their own exclusive circuit.

Consequently, we can characterize IP packet networks by saying they:

- Are based on best effort delivery.
- Require some method of error detection and assume retransmission to be the responsibility of the end stations.
- Exhibit potential for packet loss and variable delay.

IP protocols

Keep in mind the actual transmission of voice is only one aspect of digital telephony. Phone calls also require a call establishment and termination procedure using a *call control protocol*. In VoIP, there are currently three widely deployed control protocols: H.323, Session Initiation Protocol (SIP) and Skinny Client Control Protocol (SCCP).

While a detailed discussion of these protocols is beyond the scope of this paper, a brief description of each will enhance our discussion. You may remember that the purpose of these protocols is to set-up, terminate and control the calls. The protocol RTP handles the delivery of the actual voice stream.

H.323 was the first widely deployed standards-based control protocol. It was developed within the telecommunications industry and based in part on ISDN standards for the purpose of controlling multimedia calls over packet networks. As a result, it works equally well for voice or video. Most industry players, however, consider H.323 to be more complex than SIP or SCCP. Moreover, most implementations based on H.323 are only somewhat interoperable between vendors. Still, every major vendor implements at least part of their VoIP solution based on H.323.

The most significant contribution of H.323 is the use of the terms *gatekeeper* and *gateway*. The *gatekeeper*, or call manager as described by some vendors, is the device that controls all the VoIP devices within a network. It registers the phones and other devices, accepts or rejects a request to make a call, and stores the information needed to route a call, as well as a myriad of other tasks. A *gateway* is a device that moves messages between a packet environment and a non-packet environment such as the PSTN or a microwave T1 circuit. In H.323 language, a telephone is called a terminal because of the variety of other devices the protocol supports.

Many industry experts believe that SIP was developed in reaction to the complexity of H.323. In contrast to H.323, SIP was developed by the Internet community and is based on the structure of the HTTP protocol. Therefore, one could say that the vast collective experience of web developers is available to apply to telephony application development in SIP. Consequently, it uses many protocols in common with the Internet such as DNS, DHCP and ICMP. The promise of SIP is two-fold. First, endpoints will be more flexible because they can be identified by any Internet name type. Thus, host names, phone numbers, and even email accounts can be the target destination for a call. Second, systems and databases will be more interoperable because SIP is based on Lightweight Directory Access Protocol (LDAP), a protocol already in use across a wide variety of vendor platforms.

Finally, the third most popular control protocol is Cisco's SCCP, often just called *Skinny*. Skinny has proven to be simple, effective and relatively easy to troubleshoot. When a protocol analyzer captures or monitors packets from a Skinny phone that is establishing a call, the technician can see the digits dialed, the indication of ringing, the off-hook status and other useful information. This is not possible with H.323.

Table 1

Frame Number	Variation From Expected Arrival Time (in ms)
1	+4
2	-5
3	-1
4	+3
5	-2
6	+4
7	+1
8	-5
9	+2
10	-1

Network quality

Packet networks are very different from circuit-switched networks. Let us consider the three significant characteristics of a packet network, jitter, latency (or delay) and packet loss, in detail.

The first and most important characteristic of VoIP is jitter. *Jitter* is the variation in the arrival time of packets compared to the expected arrival time as measured by the receiver. For example, consider a situation where the transmitting phone is inserting two payload blocks per packet, each representing 10 ms of voice. The phone is transmitting 50 packets per second. If these packets are received uniformly every 1/50 of a second, the jitter is zero because the expected arrival times are 1/50 of a second apart and that is when the packets arrive. In other words, there is no variation from what is expected. However, suppose ten packets arrive in the manner described in Table 1.

In order to accommodate the irregularity of the arrival of packets, manufacturers use *jitter buffers*. The operation of such a buffer is illustrated in Figure 2.

Figure 2

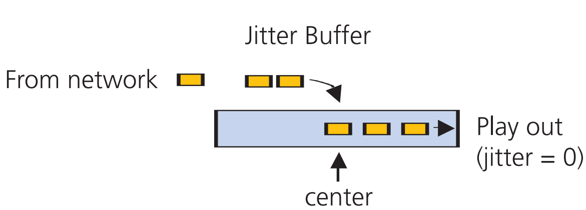

The first packet is placed in the center of the buffer. As each packet is received, it is placed behind the previous packet relative to the play out end. The size of the buffer is measured by the time it takes a packet to move from the center to the play out end and is usually configured in software. The longer the buffer, the more irregularity the buffer accommodates. The shorter the buffer, the more likely the buffer is to fill.

Two important outcomes are apparent from this operation. First, if the buffer overfills because of the early arrival of packets, they will be dropped. In a similar manner, if packets are delayed and the buffer empties, the receiver will be forced to simulate some sound or allow a period of silence. Second, the buffer introduces a delay equal to the time it takes for a packet to go from the center of the buffer to the play out end. Therefore, the longer the buffer, the more delay, known as *jitter buffer delay*, is introduced. It is therefore important to know what type of jitter your network is experiencing in order to appropriately tune your jitter buffer configuration.

How is jitter measured and reported? Generally, the receiving phone or software that is being used to monitor the connection calculates and reports the value. Each RTP packet contains a header. In that header the sender places a time stamp. Based on this time stamp and the receiver's own processor clock, the jitter variation from expected arrival time can be computed. The jitter value is computed from these jitter variations. Each receiver of audio will normally report the computed jitter value in a packet called the real time control protocol or RTCP report. Such reports are generally sent to the transmitting stations about every five seconds. Protocol Expert and Link Analyzer report these RTCP values and measure the actual jitter variation seen as the packets are monitored on a link. This is important because the reported RTCP values are often different from the jitter witnessed in different locations on the network.

Delay or *latency* is another metric that has a substantial impact on VoIP QoS. Jitter buffers will add delay to the conversation, but it is just one type of delay that needs to be considered.

Measuring transmission delay in a network can be difficult. This is because the computation of the difference between the time at which the packet is sent and when it is received must be based on a *common clock* and measurements taken at both ends of a conversation. If the sender and receiver are far apart, implementing measurement devices and connecting to such a clock may be difficult. The other way to measure this delay is through the transmission and measurement of synthetic traffic, however doing this on a production network may not be practical. For these reasons, Protocol Expert does not consider transmission delay in QoS scores. However, jitter buffer delay, codec delay, and packetization delay are calculated and used in Protocol Expert's and Link Analyzer's quality scores. Integrating these measurements is useful as they provide meaningful measures of call quality. Since transmission delay, in most typical situations, does not significantly impact quality scores it is not included.

Table 2
Types of Delay

Type of Delay	Description
Jitter Buffer Delay	Caused by the jitter buffer as it holds onto packets before sending them back out to the listener. The larger the jitter buffer, the longer the delay introduced.
Codec Delay	Caused by the codec itself and results from the fact that as the voice is sampled, it is retained briefly. This is so that the information in the current sample can be used to modify and compress the information in the next sample that is collected.
Store and Forward Delay	Caused by the action of packet relaying devices in the network such as routers, switches and firewalls. An entire packet must be stored before it can be retransmitted.
Propagation Delay	The time it takes for the bits to actually traverse the medium (wire, fiber or air)
Serialization Delay	Serialization delay is the fixed delay required to clock a voice or data frame onto the network interface. It is directly related to the clock rate on the trunk.
Packetization Delay	The time it takes an end point to place the voice payload into a packet
Transmission Delay	The combination of Store and Forward, Serialization and Propagation Delay

Finally, *packet loss* is the rate at which packets are discarded either by a device in the network or by the receiver. Protocol Expert identifies two types of packet loss. *Packet discards* refer to packets discarded by the jitter buffer. On the other hand, *packet loss* refers to packets lost in the network. Differentiating between the two is very helpful in isolating the cause of quality degradation and resolving issues. In packet networks, several things can cause packets to be dropped. First, a packet may be dropped because it becomes corrupted while traversing a link. For example, every Ethernet network interface card (NIC) checks for errors in a received frame. If it finds that a single bit is corrupted, the entire frame is dropped. So, while a link exhibits a very low probability for a single bit to be dropped, a 1250 byte frame which contains 10,000 bits has a significantly higher probability of being rejected at the receiver. In spite of this, corrupted frames in Ethernet are a minor reason for VoIP quality to be affected negatively. Frame drop rates are usually quite low.

A more significant cause of dropped frames is network congestion. When routers become congested with too much traffic, they are likely to fill their input buffers. When this happens, packets that arrive must be dropped. Routers may also drop packets as part of a quality of service scheme they are implementing. They may do this in order to provide precedence to packets that have been marked for expedited processing. Some devices in networks implement load-balancing algorithms which intentionally drop packets in order to achieve favorable levels of traffic on particular circuits.

It seems natural to ask how packet loss is detected by the receiving phone. It needs to do this because it must allow for silence during the time when the dropped packet should have been played by estimating the audio output to be played and replaying one or more previously received packets. Packets are marked by the sender with a sequence number. These sequence numbers are consecutive integers sometimes beginning with a random number and sometimes beginning with the number one. This sequence number is inserted in the RTP header of the packet. When the receiver receives a packet with an integer that is not consecutive with the previous packet, it knows immediately that a frame is missing.

Acceptable network parameters for VoIP

It has taken a number of years for manufacturers to decide on a common set of performance standards to support high-quality voice over IP. Initially, all vendors said packet loss must be less than 1%. However, as codecs became more and more sophisticated in estimating the audio to play back in the absence of a packet, they have gradually relaxed that requirement. Currently, most vendors suggest a range from 1-5%.

The acceptable level of delay is somewhat more difficult to determine. The time between when a speaker stops speaking and the speaker hears a response is sometimes called the turn-around time. Most talkers are not concerned if that time becomes as high as 500 ms, especially if they know the two parties are separated by a significant distance. This means that a one-way delay of 250 ms would seem acceptable. On the other hand, there are other factors to be considered. Echo is often aggravated by delay. Echo is a reflected copy of the original speaker's sound that comes back to the speaker's ear. It is often lower in strength and slightly delayed. In fact, echo is almost always present. Whether the speaker perceives it or not is determined primarily by the size of the reflected signal and the amount by which it is delayed. It can either be strong and delayed very little or weak and delayed significantly. In both cases, the echo should not be perceptible. But if it has both moderate delay and strength it will be perceived. Consequently, increased delay may move echo that was not perceived into the range of perceptible echo. Delay can also cause packet loss. This is because packets delivered too late to be played out when their turn comes will be lost.

As a result of such considerations, manufacturers of VoIP equipment generally have recommendations on the maximum level of acceptable packet loss, delay and jitter. Note, while these levels are defined by Cisco, Avaya and Nortel, it is not easy to define an acceptable level of quality from these metrics, or even their combination, as they do not directly translate into specific quality levels. However, using a sophisticated algorithm such as the ones provided by Protocol Expert and Link Analyzer, it is possible to generate R-factor/MOS scores that directly map to certain quality levels.

Table 3

Maximum Acceptable Levels of Packet Loss, Jitter and Delay			
	Packet Loss (%)	Jitter (ms)	Delay (ms)
Nortel[1]	5	N/A	N/A
Cisco[2]	1	30	150-200
Avaya[3]	1	20	80

1 Nortel: *Business Communications Manager 2.5: IP Telephony Configuration Guide, p.120.*
2 Cisco: *Cisco AVVID Network Infrastructure Enterprise Quality of Service Design (whitepaper), p. 1-3, Aug. 2002.*
3 Avaya: *Avaya Quality of Network Requirements, Issue 2.0, August 2002 (whitepaper).*

Voice quality

For years, the telecommunications industry used one term for acceptable telephone output, *toll quality*. However, in the last decade, people have become accustomed to levels of quality that are significantly lower. Cell phones and portable phones have contributed to this change. Today, most telecommunications professionals refer to three levels: toll quality, business quality and unacceptable.

The most widely used measure of voice quality has been the *mean opinion score (MOS)*. This score is an average rating by individuals who judge the quality on a scale from 1 (low) to 5 (high). Historically, scores above 4.0 were referred to as *toll quality*. Recently a new method was introduced using analytical techniques to calculate MOS scores. Referred to as the R-factor in Protocol Expert, it is based on packet loss, delay, jitter and an estimate of user perception. It has been found to correlate well to the more traditional MOS score. Fluke Networks' Optiview Protocol Expert makes two R-factor measurements. First the value which depends on jitter and packet loss is computed and reported as the *Network R-factor*. Then Protocol Expert adds the user perception value and reports the *User R-factor*. This measure adds consideration for the *recency* of the impairment. Research shows that negative characteristics in a call have less affect if they happen in the past rather than near to the time they are heard. As mentioned above, both R-factors correlate well to MOS scores. Table 4 shows the approximate relationship.

Table 4

Desirability of Quality	R-Factor	MOS
Desirable (Toll quality)	80-94	4.4 – 5.0
Acceptable (Business quality)	70-80	3.6 - 4.4
Reach Connection (acceptable)	50-70	2.6 - 3.6
Unacceptable	0-50	0 - 2.6

The quality of voice is a subjective judgment to users. While the recommendations for network characteristics are given above, experiments have revealed two important facts. First, VoIP phones and gateways have improved the way they deal with jitter, loss and delay significantly over the last few years. Second, poor quality output is more evident when a *combination* of jitter, loss and delay are present in the network. For example, 35 ms of jitter created only a slight decrease in voice quality as long as the packet loss and delay were insignificant. OptiView Protocol Expert and Link Analyzer use a sophisticated algorithm based on the E-Model (ITU G.107) to generate R-factors and estimated MOS scores. These objective measurements can be used to judge a VoIP deployment's relative quality, as well as monitor changes to quality over time, without the reliance on subjective user measurements. Three different MOS quality scores are provided since different algorithms yield slightly different scores.

- The MOS-LQ (listener quality) scores map the combined R-factors to a Mean Opinion Score for listener quality. The MOS-LQ estimate does not contemplate the impairments that cause conversational quality problems, such as delay, and can be compared to subjective MOS scores.
- The MOS-CQ (call quality) score estimates conversation call quality expressed as a MOS score. The estimated MOS-CQ maps the combined R-factors to a Mean Opinion Score for conversational quality. MOS-CQ incorporates factors, such as delay and recency, which affect conversational quality, and implicitly affect listening quality.
- The MOS-PQ (PESQ) score estimates PESQ call quality expressed as a MOS score. The MOS-PQ maps the R-factors to the ITU-T standard P.862 (PESQ) end-to-end measure of voice quality using a test signal.

While other solutions provide a single MOS score, Protocol Expert and Link Analyzer let you compare the quality of VoIP calls based on several different algorithms.

Measurements with Fluke Networks' Optiview Protocol Expert

Optiview Protocol Expert and Link Analyzer are excellent tools for assessing a network's readiness to support VoIP. They can also discover problems after VoIP has been deployed. The VoIP feature option for Protocol Expert provides a rich set of facts and measurements related to RTP and call set-up of VoIP calls. A Call Detail record lists sender and receiver identifiers (which may be customized as IP address, phone number or other identifiers), statistics, quality (which may be customized as various MOS scores, R-factors or jitter) for both directions and the associated call control protocol. Key metrics such as jitter and packet loss are measured and can be compared to the statistics indicated within RTCP. All of this can be displayed in graphical form or exported to a database with the click of a few buttons.

Figure 3 shows the Optiview Protocol Expert VoIP Properties screen. Everything that is critical to assessing the quality of the call is here: jitter, packet loss, estimated packets discarded, user and network R-factor. Drill down analysis of each individual call is accessible by double clicking on the call you are interested in as shown in Figure 4.

Figure 4
Detail Call View

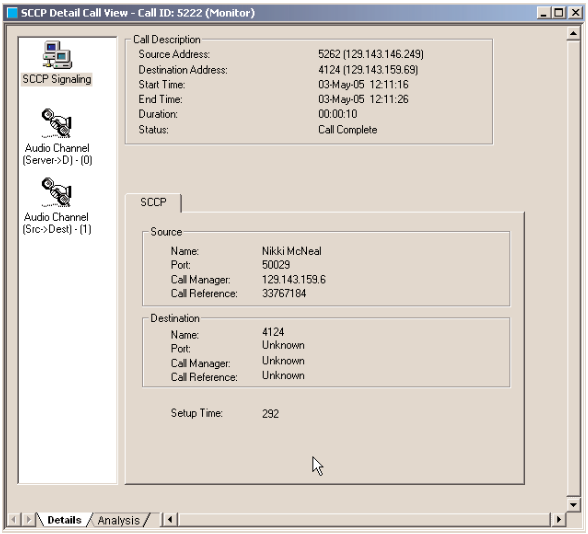

Extensive information is available from the Channel Detail view. Highlighting the signaling data icon on the left shows the call source and destination names, IP address, port number, call manager and call reference number. This information is especially helpful because troubleshooting issues can be linked back to the call manager. Simply clicking on either of the Channel Detail icons provides detailed quality-related information about

Figure 3
VoIP Properties

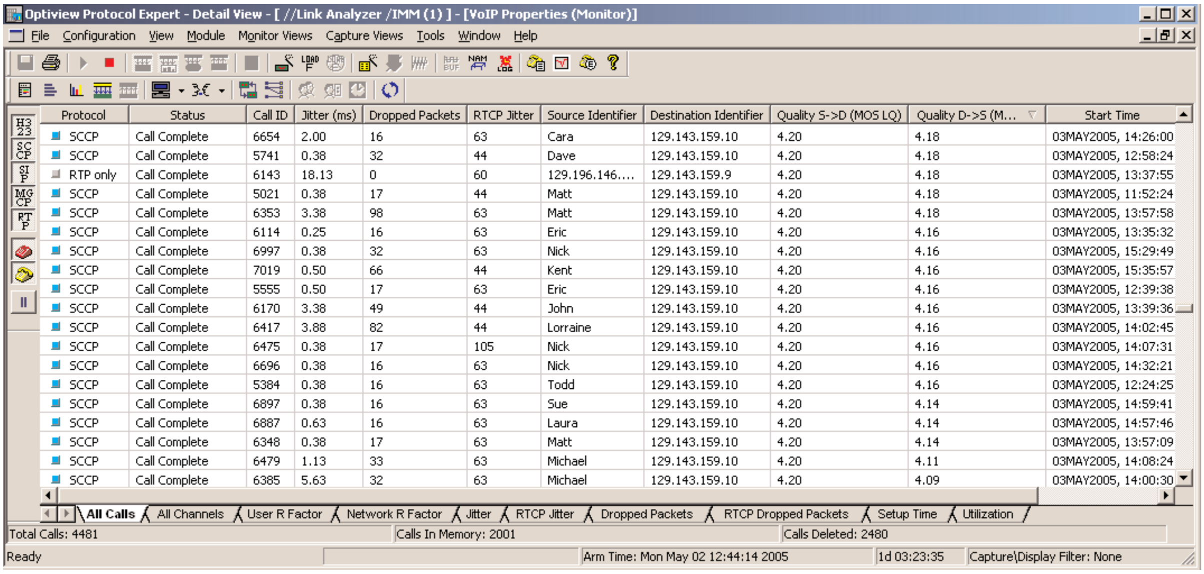

Figure 5

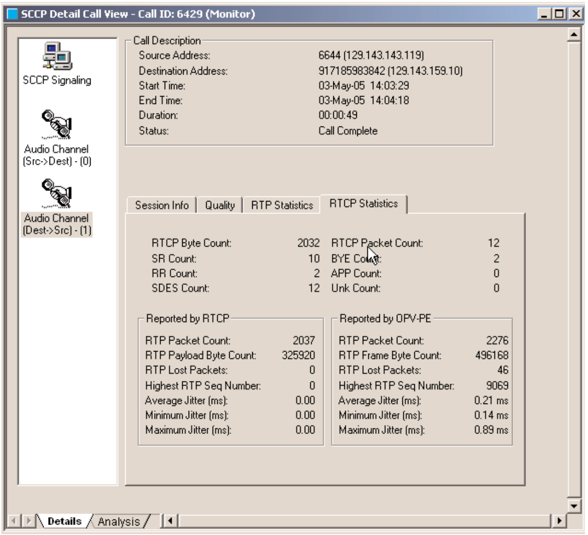

Figure 6

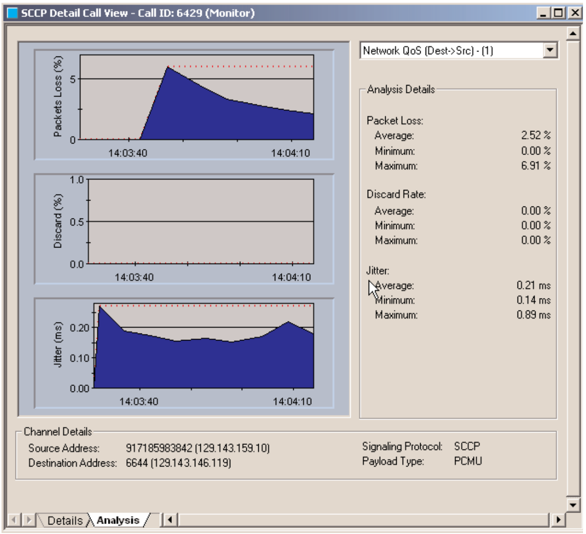

the individual call as shown in Figure 5. By clicking on the RTP Statistics tab, you will get detailed information that has been calculated from the RTP packets received. This will include the port numbers, bytes transferred, packets lost and jitter.

In Figure 6, a complete summary is created when you click on the Analysis tab at the bottom of the window. This is helpful in assessing the performance of the channels across time and in identifying network problems.

Now that we have seen how Protocol Expert can be used to reveal measurements about call features, let's review the results from some actual implementations. Figure 7 below shows a call with less than optimal network characteristics.

Opening the Call Detail shows the phones involved in the voice exchange and the codec type being used, such as G.711, as indicated by the PCMU (pulse code modulation, mu-law) designation. Looking at the quality scores for each individual channel shows a problem in the network in one direction of flow: 156 packets were dropped. As a result, the quality MOS score is marginal. This prompts us to look deeper. In the analysis tab of the Call Detail window, we

Figure 7
VoIP Properties

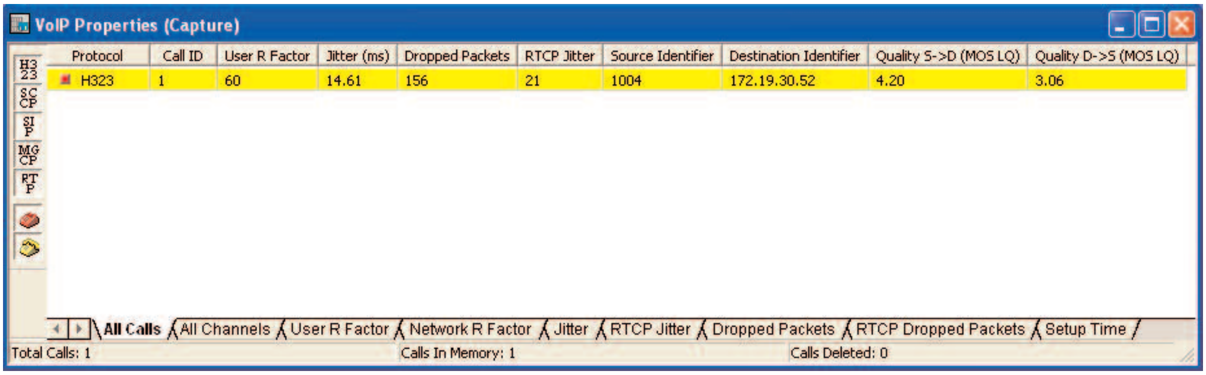

Figure 8
VoIP Call Detail Analysis

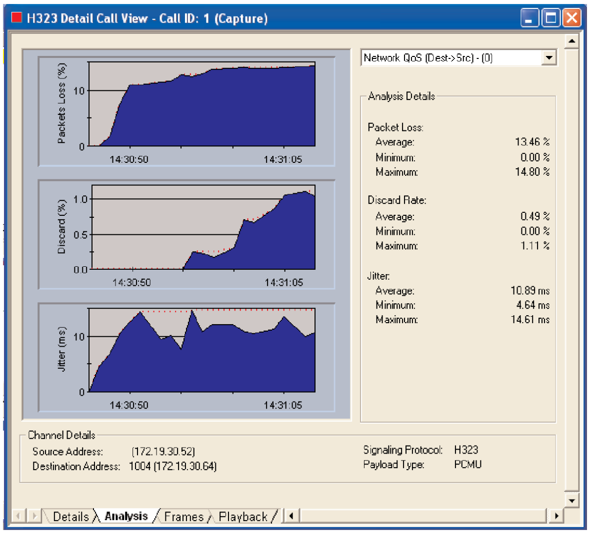

Figure 9
VoIP Call Detail Analysis of the Reverse Channel

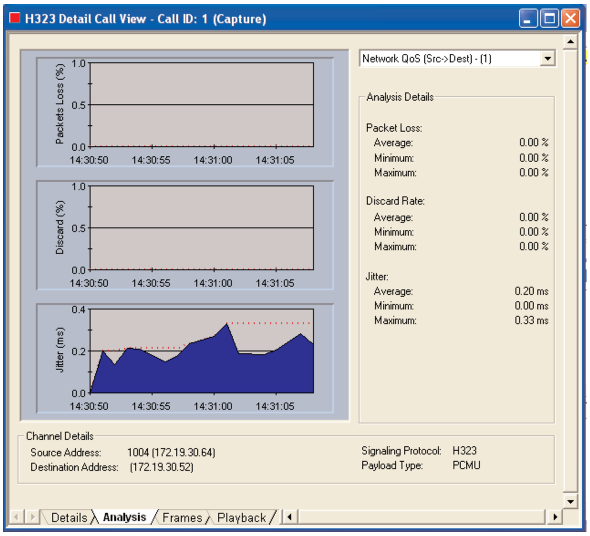

see the screen in Figure 8. It reveals the fact that all of the packet loss appears to be occurring in a single channel (destination to source).

By clicking the drop down menu and selecting the opposite channel (source to destination), we can verify that the degradation is only occurring in one direction. This is shown by the screen in Figure 9 which indeed shows no packet loss in the opposite direction.

Protocol Expert shows the type of detailed information that is necessary to troubleshoot the typical symptom provided by end users – "call quality is bad." Notice also that the jitter value in the problem channel reported by RTCP is not exactly the same as that measured by Protocol Expert.

This is indicative of end point management reports that have no way to identify network degradation apart from end user experience at the phones. In order to obtain an accurate picture of call quality, remote link monitoring and analysis must be done with a solution like Link Analyzer and Protocol Expert. Further investigation of this call revealed that a defective patch cable attached to a switch in the path of the call had a defective wire and was causing significant cross-talk on the channel. OptiView Protocol Expert and Link Analyzer provide a means of isolating call quality problems to locations on the network and individual devices. This speeds problem resolution and makes for happier end users.

Using Protocol Expert to compare codecs

There has been a lot of discussion regarding the question of whether the selection of a codec can affect the quality of voice output in different situations. Protocol Expert can help you make the decision about which codec you want to use.

Consider Figure 10. Five calls were under way when Protocol Expert was started. We can tell the calls were already started since the call set-up protocol was not recorded. As a result Protocol Expert reports the call as an RTP stream only.

Changing the codec can impact the quality of the voice as we shall see. Figure 11 shows part of the decode screen for the five calls we have been analyzing. There are several ways to see that the codec being used is G.729. First, the packet size is 78 bytes. This is the normal default when 58 bytes of Ethernet-IP-UDP-RTP headers precede two 10-byte blocks of G.729 encoded voice. Second, by opening the RTP header field in the detailed portion of the screen, we can see that the payload is labeled as G.729 by Protocol Expert. This information can also be obtained by looking at the "All Channels" tab which identifies the codec for each channel.

Figure 11
Decoded Calls

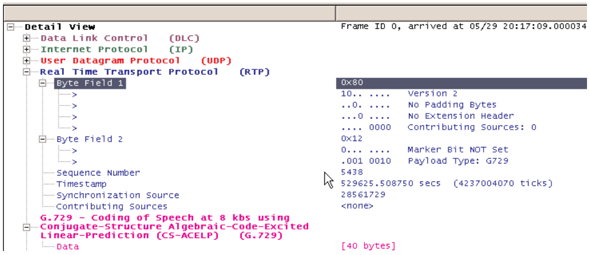

Figure 10

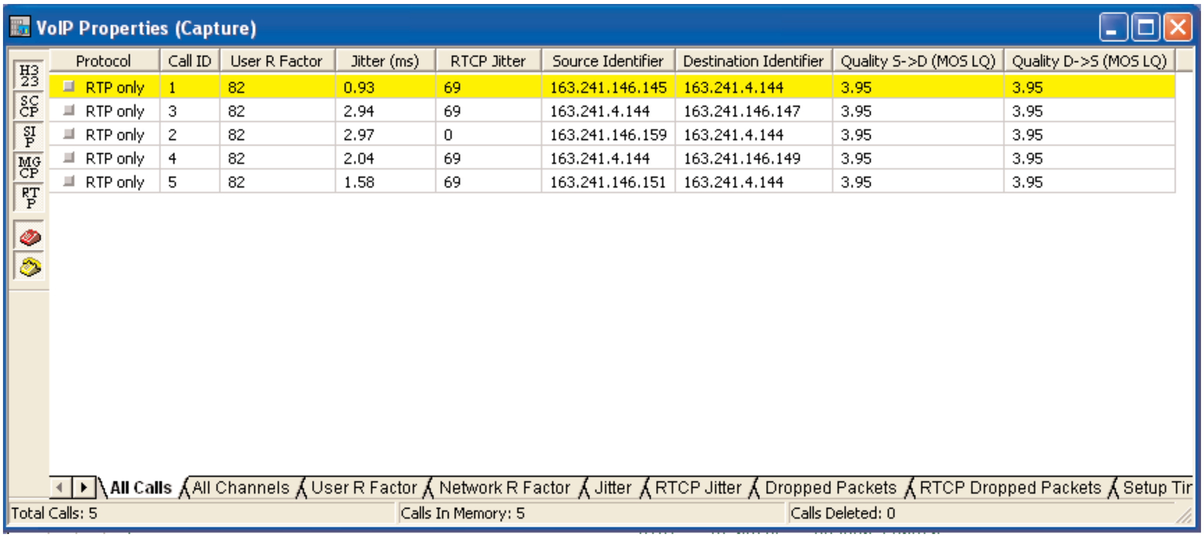

Figure 12
Jitter on Five Calls with G.711 Codec

VoIP Properties (Capture)

Protocol	Call ID	User R Factor	Jitter (ms)	RTCP Jitter	Source Identifier	Destination Identifier	Quality S->D (MOS LQ)	Quality D->S (MOS LQ)
RTP only	5	92	1.10	64	163.241.146.145	163.241.4.144	4.20	4.20
RTP only	3	92	3.95	64	163.241.4.144	163.241.146.159	4.20	4.20
RTP only	1	92	0.31	64	163.241.4.144	163.241.146.147	4.20	4.20
RTP only	4	92	0.22	64	163.241.146.149	163.241.4.144	4.20	4.20
RTP only	2	92	1.05	64	163.241.4.144	163.241.146.151	4.20	4.20

All Calls / All Channels / User R Factor / Network R Factor / Jitter / RTCP Jitter / Dropped Packets / RTCP Dropped Packets / Setup Time

Total Calls: 5 Calls In Memory: 5 Calls Deleted: 0

By changing the codec setting within the VoIP equipment to G.711 and observing the same five calls again, Figure 12 reveals a significant improvement in quality.

By comparing Figure 10 with Figure 12, we can see that both the User R-factor and Network R-factor have increased significantly from business quality to toll quality. This has occurred in spite of only a slight improvement in the RTCP jitter reported by the phones.

Figure 13
IP Phone Report

13:31:01 01/25/05	7175551107
Call Statistics	
RxType:G711	TxType:G.711
RxSize: 20 ms	TxSize: 20 ms
RxCnt: 001573	TxCount: 001520
AvgJtr:16	MaxJtr : 20
RxDisc: 0000	RxLost: 0000
Press "i" to cancel...	
Exit	

Qualifying the statistics reported by the phone

Today most IP phones provide a report of jitter and packet loss on the display along with other important parameters such as the phone's IP address. Figure 13 shows a replica of the screen that is provided by a Cisco 7960 phone.

In comparing these values to those measured by Protocol Expert, we can see that there is a discrepancy between what the phone reports and what is reported by Protocol Expert (.18 ms vs. 20 ms). This indicates that a jitter issue is only occurring in one direction and is somewhere between the Link Analyzer and the phone. Figure 14 shows the Link Analyzer measurements of the same call and verifies the jitter values at the Link Analyzer's location.

Figure 14
Comparison to Phone Report

VoIP Properties (Capture)

Channel	Call ID	Source Identifier	Destination Identifier	Last Seen	CODEC	Max Jitter (ms)	Quality (MOS LQ)	RTCP...
Audio	6	1101	1107	25JAN2005, 13:21:53	PCMU	0.20	4.20	N/A
Audio	6	1107	1101	25JAN2005, 13:21:53	PCMU	0.18	4.20	N/A
Audio	5	1100	1104	25JAN2005, 13:21:53	PCMU	0.17	4.20	N/A
Audio	5	1104	1100	25JAN2005, 13:21:53	PCMU	0.16	4.20	N/A

All Calls / All Channels / User R Factor / Network R Factor / Jitter / RTCP Jitter / Dropped Packets / RTCP Dropped Packets / Setup Time

Total Calls: 2 Calls In Memory: 2 Calls Deleted: 0

Pinpointing a QoS problem with OptiView Protocol Expert

OptiView Protocol Expert is an excellent tool for finding the source of a QoS problem such as degradation in quality. We will illustrate this with an example involving a VPN connection.

Suppose a company has a VPN tunnel supplied by a carrier between its headquarters and a remote office, as shown in Figure 15. In addition, suppose the quality of calls between X2001 and X3001 has degraded since the introduction of the VPN tunnel. In order to separate the impact of the VPN from the impact of the separate corporate networks, we take jitter and loss measurements using Protocol Expert. Then, we ping to estimate delay and record all of this in Table 5.

We can see from the table that as we measure at different points on the network, the packet loss and jitter measurements observed help us isolate the area of degradation. When measuring the channels close to the phones traveling to the VPN we see good scores, which means that the local networks are not introducing any problems. However when we watch the channel details coming out of the VPN, we see a substantial reduction in the quality of the measurements. Jitter jumps by 15ms (17 minus 2) for the 2001➡3001 channel, and 16ms for the 2001 ⬅ 3001 channel. Likewise, the additional delay caused by the VPN carrier is 20ms (24 minus 4) as observed from the ping measurement. This effect is quantified in the channel information observed in User R-factor when measuring on either side of the VPN tunnel.

Table 5
Measurements Across the Network

Parameter		Point of Connection			
	Channel	A	B	C	D
Packet Loss	2001➡3001	0%	0%	3%	3%
	2001⬅3001	2.7%	2.7%	0.1%	0.0%
RTCP Packet Loss	2001➡3001	3%	3%	3%	3%
	2001⬅3001	2.7%	2.7%	2.7%	2.7%
Jitter	2001➡ 3001	1 ms.	2 ms.	17 ms.	18 ms.
	2001⬅3001	20 ms.	19 ms.	3 ms.	3 ms.
RTCP Jitter	2001➡3001	18 ms.	18 ms.	18 ms.	18 ms.
	2001⬅3001	20 ms.	20 ms.	20 ms.	20 ms.
Ping Delay to X3001		0 ms.	4 ms.	24 ms.	27 ms.

We also noticed that the RTCP jitter and loss remain about the same across the network. This is to be expected because these values are measured by the endpoints, the phones in this case. But more importantly, this points out the added value in using Protocol Expert and Link Analyzer on the network – by helping us track a problem to its source despite the fact that all endpoints were experiencing the same problem. While the phones tell us there is a problem, they do not help us *find* the problem. Protocol Expert and Link Analyzer allow us to pinpoint the location of the problem to the VPN tunnel.

Figure 15
VoIP Over A VPN Connection

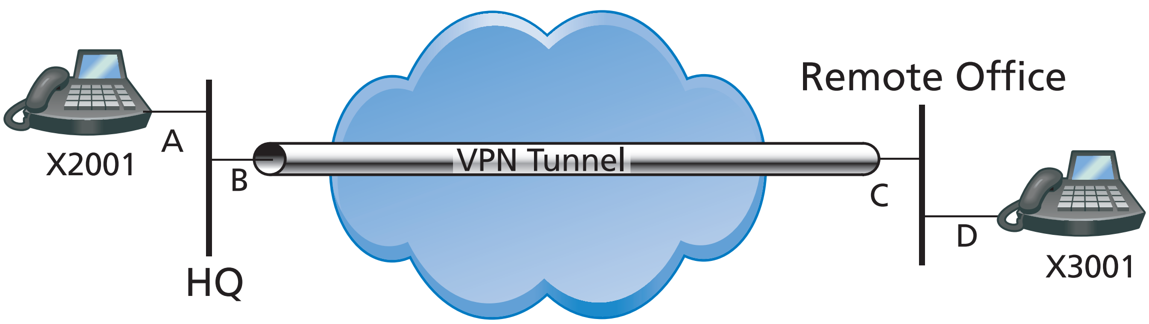

Summary

In this paper, we have considered how voice is converted into packets and transmitted on IP networks. We have investigated the primary characteristics of networks that can impair the quality of voice: jitter, delay and packet loss. We have examined how Fluke Networks' Optiview Protocol Expert can be used to measure these values. It can assist us in assessing the quality of the voice output, either as a MOS score or an R-factor. From our study, we were able to determine that R-factors will be affected by the codec selected as well as the combination of values of the network parameters. We also learned the point at which the network parameters affect the values. Jitter measured close to the device sending the RTCP reports will be quite different from the values measured near the receiver of the reports. We also studied how Protocol Expert can be used to pinpoint the source of a problem, something which is often impossible based solely on the information obtained from the endpoints.

Protocol Expert is a valuable tool for understanding, troubleshooting, monitoring and managing VoIP networks. It provides a broad range of critical measurements obtained in real time on the network or post packet capture using a trace file. Protocol Expert supplies extensive detail about the network and call characteristics that determine the quality of service VoIP end users experience. By monitoring these characteristics, network engineers can be alerted to degradation as it happens. They can also effectively troubleshoot the problem and isolate the root cause much faster and easier than with embedded metrics.

For VoIP end point analysis and troubleshooting during moves, adds and changes, be sure to check out Fluke Networks NetTool for VoIP. This small handheld device lets you analyze call set up and performance for individual end users.

Printed in U.S.A. 6/2005 2279730 A-US-N Rev A