

Voice over Cable (VoCable)

White Paper

Edward Morgan
Executive Director, R&D

Telogy Networks, Inc.
A Texas Instruments Company

Debbie Greenstreet
Senior Product Manager

Telogy Networks, Inc.
A Texas Instruments Company

February 2000
Version 1.6
SPEY001

I. Introduction

Cable-based IP telephony holds the promise of simplified and consolidated communication services provided by a single carrier at a lower total cost than consumers currently pay to separate Internet, television and telephony service providers. Cable operators have already worked through the technical challenges of providing Internet service and are optimizing the existing bandwidth in their cable plants to deliver high-speed Internet access. Now, cable operators have turned their efforts to the delivery of integrated Internet and voice service using that same cable spectrum.

Cable-based IP telephony falls under the broad umbrella of voice over packet (VoP), meaning that many of the challenges facing cable operators are the same challenges that telecom carriers face as they work to deliver voice over ATM and Frame Relay networks. However, ATM and Frame Relay services are targeted primarily at the enterprise, a decision driven by economics and the need for service providers to recoup their initial investments in a reasonable amount of time. Cable, on the other hand, is targeted primarily at the home. Unlike most businesses, the overwhelming majority of homes in the United States are passed by cable, reducing the required up-front infrastructure investment significantly.

While cable isn't without competition in the consumer market -- xDSL has emerged as the leading alternative to broadband cable -- cable operators are positioned well to capitalize on the convergence trend if they are able to overcome the remaining technical hurdles and deliver telephony service that is comparable to the public switched telephone system. This paper discusses the challenges of providing cable-based IP telephony and offers an overview of Texas Instruments' unique software architecture designed to enable clear, reliable IP-voice solutions. Texas Instruments' Telogy Software™ products are key components in TI's comprehensive solutions that also include programmable DSPs and a complete cable modem chip set, significantly reducing development costs and accelerating time to market for equipment providers.

II. Making Cable a Viable Telephony Medium

For almost a century, Americans have taken for granted the nearly unailing service provided by the public telephone network, often referred to as plain old telephone service (POTS). If cable is to emerge as a legitimate alternative, there are many technical issues that must be addressed. Perhaps the most fundamental of these is the evolution of the nation's cable infrastructure from a one-way, broadcast medium to a two-way, personal-communications medium.

Half Duplex v. Full Duplex Cable Infrastructure

Cable was first introduced in the U.S. in the late 1950s. For the next 30 years, nearly every mile of buried cable was half duplex. This means that it was capable of broadband transmission in the downstream direction, i.e. from the head end to the subscriber, but not in the upstream direction. Communication from the subscriber back to the head end was possible only via a telephone line.

This makes half-duplex lines cumbersome even for premium TV services, such as pay-per-view, that require upstream communication. It makes half-duplex lines extremely inconvenient for Internet service due to the fact that outbound email messages and HTTP requests have to be sent via the phone. And, it renders half-duplex lines completely useless for voice because such service requires packets to be sent up and down stream continuously.

In recent years, cable operators have been investing heavily to upgrade their buried cable from half to full duplex as a necessary first step to capitalize on the demand for integrated data and voice services. While upstream transmissions still are not as fast as downstream (typically 1.5 to 3 Mbps downstream and 500 Kbps to 2.5 Mbps upstream), full-duplex lines offer sufficient throughput to support cable-based IP telephony. As cable operators compete for subscribers with xDSL providers, the speed with which cable operators replace older lines with full-duplex lines will be critical to their ultimate success.

Telephony Service across a Broadcast Media

Unlike POTS, which was developed from the outset as a point-to-point communication technology, cable networks were designed originally to broadcast one signal to many recipients. There was no concept of dedicated circuits and there was no need to parcel out bandwidth to individual subscribers. To enable cable-based IP telephony, modifications must be made to the way bandwidth is allocated and packets are delivered. This must be done without using the bulk of the cable spectrum because most of the bandwidth will continue to be used for TV broadcasts.

Direct Connect. Callers must be able to send and receive only their own voice packets, and these packets must be given priority over data packets to ensure that callers experience smooth, uninterrupted conversations. The first step in this process was addressed by the Data over Cable Service Interface Specification (DOCSIS). DOCSIS established universal ground rules for the transmission of packets across cable networks, ensuring that packets won't be routed incorrectly.

DOCSIS was later enhanced (in version 1.1) with quality of service (QoS) and security features that are necessary for voice communication. DOCSIS 1.1 also enables the prioritization of packet traffic. This allows cable operators to give certain packets (i.e. voice) the right of way and allows other traffic to be sent with a "best effort" priority as determined by bandwidth availability. However, even this second-generation DOCSIS standard was not intended to address all of the technical issues associated with cable-based voice service.

To fill in the gaps left by DOCSIS, CableLabs[®]¹ created the Packet Cable specification known as the Network-based Call Signaling (NCS) protocol for signaling voice calls over cable networks. NCS leverages the existing Media Gateway Control Protocol (MGCP) and the protocol is sometimes referred to as MGCP NCS. NCS uses network-based *call agents* to negotiate cable-based IP telephony calls. Call agents, which will be discussed later in this paper, ensure that voice packets traverse the network and are audible only at the two conversation end points, i.e. the people on either end of the telephone.

Security. While POTS is considered an extremely secure service, cable-based IP telephony is not. Much like cellular telephony, cable-based conversations are susceptible to illegal wire tapping and inadvertent "chat" conditions. To address this untenable situation, DOCSIS and NCS support multiple security services

NCS currently supports the IPsec authentication specification. Adequate protection of telephony connections can be achieved if the telephony gateway accepts only packets that have been authenticated by IPsec. As described in the CableLabs online magazine *Specs Technology*, DOCSIS supports an encapsulation protocol for encrypting packet data across the cable network. The encapsulation protocol defines the frame format for carrying encrypted packet data, the set of supported data-encryption and authentication algorithms, and rules for applying the cryptographic algorithms to packet data.

CableLabs further reports that DOCSIS currently employs the Cipher Block Chaining (CBC) mode of the U.S. Data Encryption Standard (DES) to encrypt packet data. The protocols are extensible, can support multiple encryption algorithms and will, in all likelihood, be extended to support the new Advanced Encryption Standard (AES) once it is in place.

Power Consumption

As most people know, traditional telephones draw all the power they need from POTS lines. Because the public phone system has evolved to such a reliable state, and it is essentially immune from the effects of power outages, it is exceptionally rare that service is lost. Electrical utilities in most areas do *not* offer this degree of unfailing reliability. Therefore, head-end and customer-premises cable equipment that relies solely on the local electric company for power puts users at risk of losing phone service should a power outage occur.

To address this issue, “lifeline service” requirements are being implemented across the country that require IP phones, such as those that connect to cable lines, to provide at least four hours of battery back up. To meet this requirement, equipment manufacturers must develop phones that can be powered by as little as three watts. A key to achieving this is a telephony chipset that minimizes idle processing cycles and offers sufficient onboard memory to handle all signal processing.

The ideal cable-based IP telephony system is typically built with a RISC microprocessor to handle the signaling functions and digit collection. The necessary telephony peripherals, such as a LAN controller and USB, are on a single chip to conserve power and dedicated hardware should be used for the cable-communications protocol. Several megabytes of high-speed RAM are needed for the signal processing and the same amount of non-volatile memory is needed to store the telephony application. The non-volatile memory should be electrically re-programmable, such as a FLASH memory, to enable online software updates.

A high-performance, low-power DSP is needed to support the analog functionality, e.g. codecs, noise reduction and echo cancellation. A programmable DSP can greatly reduce application-development time for solution providers. TI’s TMS320C54x DSP is one such chip.

Billing

Telephony billing is an extremely complex process. Most cable TV customers receive the same bill each month. Aside from pay-per-view requests, there is no need to meter or monitor customer usage. Telephone billing is quite different. A typical bill includes recurring monthly service fees, international and long-distance charges that vary based on time and day, and premium services, such as *69 and directory assistance, that are billed on a per-use basis.

To enable timely, accurate billing, call agents or network interfaces (NI) must collect all relevant usage data. The NI is the cable equivalent to the phone box that is outside every home. In the absence of a NI, cable-based IP telephony can also be delivered using voice-enabled cable modems inside a customer’s home. If the call agents collect the billing data, the NI or cable modem do not have to be involved. Otherwise, the software inside the NI or cable modem must provide application programming interfaces (APIs) so that the billing system can access the relevant data. Depending on each cable operator’s implementation, the data may be contained in standard management information base (MIB) files or in unique files set up specifically for telephony metering.

III. Challenges of Providing Toll-Quality Telephony Service on an IP Network

For cable operators, choosing which standard to support and preparing their infrastructures to support voice is only the beginning of the technological obstacle course. What remains is the quality of service (QoS) challenge inherent in all VoP implementations. Among the most significant QoS hurdles are transmission latency, echo, jitter and lost packets. These QoS factors are relatively harmless for data transmissions but must be dealt with aggressively to provide acceptable voice quality.

Latency

Latency, or delay, causes two problems -- echo and talker overlap. Echo is caused by the signal reflections of the speaker’s voice. Echo becomes a significant problem when delay is greater than 50 milliseconds. Since echo is a significant quality problem, equipment providers must implement echo cancellation. Talker overlap becomes significant if one-way delay is greater than 250 milliseconds, so every effort must be made to minimize delay. The sources of delay in a VoP implementation include:

Accumulation Delay. This delay is caused by the need to collect a frame of voice samples to be processed by the voice coder. It varies from a single sample time (.125 microseconds) to many milliseconds. Standard voice coders (and their frame times) include:

- G.728 LD-CELP -- 2.5 milliseconds
- G.729a, b, e CS-ACELP -- 10 milliseconds

Algorithmic Delay (sometime called look-ahead delay). This delay is caused by the characteristics of a specific voice encoding algorithm. An example of algorithmic delay is:

- G.729a, b, e CS-ACELP – 5 milliseconds

Processing Delay. This delay is caused by the actual process of encoding and collecting samples into a packet for transmission. The encoding delay is a function of both the processor execution time and the type of algorithm used. Often, multiple voice-coder frames will be collected in a single packet to reduce overhead. For example, three frames of G.729 codewords, equaling 30 milliseconds of speech, may be collected and packed into a single packet. This process of encapsulating several small packets into a single larger frame is called concatenation.

Network Delay. Network delay is a function of the processing that occurs as packets are sent across a network. This delay is caused by the physical medium and the protocols used to transmit the voice data, and by the buffers used to remove packet jitter on the receive side. The jitter buffers add additional delay that is used to smooth the jitter created by the varying times at which each packet arrives. This delay can be a significant part of the overall delay since it can be as high as 70-100 milliseconds.

Polling Delay. Cable-based IP telephony creates an additional latency that other packet networks do not because of the way head-end systems collect packets from callers. The head end polls the NI at each customer location. Because the head end doesn't maintain a continuous connection with each NI, there is a transmission delay while voice packets wait for the next poll. Therefore, it is important that cable-based IP telephony equipment minimize this delay by anticipating when the next poll will arrive -- a process called grant synchronization -- so that the packets are queued up and ready to go.

Echo

Echo is present even in a conventional POTS network. However, it is acceptable because delay is less than 50 milliseconds and the echo is masked by the normal side tone that every telephone generates. Echo becomes a problem in VoP networks because the delay is almost always greater than 50 milliseconds. Thus, echo-cancellation techniques must be used. The International Telecommunication Union (ITU) standards G.165 and G.168 define performance requirements for echo cancellers.

Echo is generated toward the packet network from the telephone network. The echo canceller compares the voice data received from the packet network with voice data being transmitted to the packet network. The echo from the telephone network is removed by a digital filter on the transmit path into the packet network.

Jitter

The delay problem is compounded by the need to remove jitter -- a variable inter-packet timing caused by the fact that packets do not all cross the network at the same speed. Removing jitter requires collecting packets and holding them long enough to allow the slowest packets to arrive and be played in the correct sequence. This causes significant delay. The conflicting goals of minimizing delay and removing jitter have led to various schemes aimed at optimizing the size of the jitter buffer to minimize its impact on latency.

A common approach in cable-based IP telephony is to count the number of packets that arrive late and create a ratio of these packets to the number of packets that are successfully processed. This ratio is then used to adjust the jitter buffer to target a specific late-packet ratio.

Lost Packets

In today's IP networks, voice frames are treated exactly like data. Under peak loads and congestion, voice frames will be dropped at the same rate as data frames. The data frames, however, are not time sensitive and dropped packets can be corrected through retransmission. Lost voice packets cannot be handled in the same manner. Some techniques used by VoP software to address this problem:

Interpolation. Interpolate for lost speech packets by replaying the last packet received during the interval when the lost packet was supposed to be played out. This scheme is a simple method that fills the time between non-contiguous speech frames. It works well when the incidence of lost frames is infrequent. It does not work very well if there are a number of consecutive lost packets or a burst of lost packets.

Redundancy. Send redundant information at the expense of bandwidth utilization. The basic approach replicates and sends the Nth packet along with the (N+1)th packet. This method has the advantage of being able to correct for the lost packet exactly. However, this approach uses more bandwidth and also creates greater delay.

Voice Coder. A hybrid approach uses a lower-bandwidth voice coder to provide redundant information carried along in the (N+1)th packet. This reduces the problem of bandwidth consumption, but does not solve the problem of delay.

IV. Fax Offers Additional QoS Challenges for Cable Networks

The challenges of implementing fax over cable networks are similar to those of voice. The two most significant issues are timing and lost packets. The delay of fax packets through a packet network causes the precise timing that is required for the fax protocol to be skewed and can result in the loss of the call. The Fax over Packet protocol compensates for the skewed timing of messages so calls are not dropped and the accuracy of faxed images is not compromised.

Lost packets can be an even more serious problem for IP-fax systems than for IP-voice systems. In a VoP application, the loss of packets can be addressed by replaying last packets and other methods of interpolation. Fax over Packet applications, however, have more severe constraints on the loss of data because the fax protocol can fail if information is lost. The severity of the problem varies depending on the type of fax machine and whether Error Correction Mode is enabled.

CableLabs is working on specific fax services that will be added to NCS to standardize the implementation of fax over cable networks. There is currently an optional fax relay service in the NCS protocol.

V. International Standards

Recognizing the importance of standards and compatibility in helping to market new services, and because of the slow progress on IEEE 802.14, North American cable system operators have settled firmly on the DOCSIS/Multimedia Cable Network System Partners (MCNS) cable modem specification developed by CableLabs, their own research and development group. The International Telecommunication Union (ITU) has also accepted DOCSIS, calling it ITU J.112. Even an accepted standard continues to evolve. In April 1999, CableLabs issued DOCSIS 1.1, adding support for IP telephony and other constant-bit-rate services. Since the newer standard is backward compatible, DOCSIS 1.0 and 1.1 cable modems can operate in the same spectrum on the same network. CableLabs is already at work on the DOCSIS 1.2 specification, which includes next generation physical layer technology intended to enable higher speed two-way services and allow cable companies to increase their networks' data capacity.

The situation is quite different in Europe, where two large consortia are vying to win acceptance of their competing standards.


One group, the European Cable Modem Consortium, composed of manufacturers seeking to leverage their North American MCNS/DOCSIS experience and investment, came together in October 1998 to develop and promote an ITU-compliant standard for cable modems, set-top boxes and head-end equipment. Members include Broadcom, Cisco, Dassault-AT, Deltakabel, Elsa, FUBA/General Instrument, Motorola, Pace, Samsung, Teldat, Tonna and 3COM. "EuroDOCSIS," as their standard is known, is DOCSIS with the addition of an 8MHz bandwidth downstream channel (within a 100 to 860 MHz spectrum), ITU-T J.83 A Forward Error Correction and upstream bandwidth of 5 to 65 MHz. These adaptations are all accomplished in the physical layer, leaving the media access control layer unchanged.

These modifications make EuroDOCSIS compatible with its major European competitor: Digital Video Broadcasting, developed by the DVB/DAVIC (Digital Audio Video Council) Interoperability Consortium. Coincidentally or not, this consortium was also founded in October 1998 by Alcatel, Cocom, DiviCom, Hughes Network Systems, Nokia, Sagem, Simac, Thomson Broadcast Systems and Thomson Multimedia. Though the European standards battle is not necessarily over, DVB appears to be winning. In April 2000, the European Telecommunications Standards Institute (ETSI) officially accepted DVB for the interaction channel for cable systems. The DVB/DAVIC Interoperability Consortium claims this makes DVB the only accepted standard in Europe for data communication over HFC networks. ETSI ES 200 800, as the standard is known, covers both cable modems and set top boxes and meets the requirements established by the European Cable Communications Association (ECCA) specifications for the “EuroModem” and “EuroBox.” Cable operators belonging to the EuroModem Consortium, with over 25 million subscribers among them, have estimated an immediate need for up to 500,000 EuroModem cable modems as soon they become available. The first mass roll-out was announced in August 2000, when industree BV, of the Netherlands signed an agreement to supply a minimum of 20,000 units in 12 months to Essent Kabelcom, one of the Netherlands largest cable operators.

VI. Embedded Software Architecture

Texas Instruments has developed Telogy Software™ products, embedded VoP software for cable-based IP telephony. The software supports cable modems and NIs (up to four ports) as well as the telephony gateway (up to several thousand ports) at the cable head end. The software supports the MGCP protocol discussed earlier in this paper as well as the Session Initiation Protocol (SIP). The basic purpose of the two protocols, i.e. to process packetized voice traffic, is the same. However, TI's Telogy Software™ products support both because standards bodies are divided as to the relative merits of each.

Figure 1: Voice over Packet Software Architecture



As shown in Figure 1, The Telogy software interfaces to both streams of information from the telephony network and converts them to a single stream of packets transmitted to the packet network.

MGCP is a centralized call-processing system in which the intelligence resides primarily at the head end. The cable modems and NIs are similar to “dumb” clients and the system relies on call agents to negotiate the call through the network. SIP is a distributed system in which the intelligence resides in the NI and the head end is mainly a gateway to the public telephone network.

The greatest benefit of MGCP implementations is simplified, efficient management and administration. Fault detection and isolation are typically limited to the head end. And, there is no need to distribute software upgrades and patches to customers, meaning that there is also no concern about software-version synchronization among all NIs.

The supporters of SIP, on the other hand, argue that it is a more scalable and reliable system. The case for scalability is due to the fact that the head end, acting mainly as a gateway, is unlikely to bottleneck subscriber capacity. Traffic load would be the only potential bottleneck, not processing time. Supporters also claim that SIP is more reliable because a SIP-based network architecture does not have a single point of failure.

TI's Telyogy Software™ products MGCP- and SIP-compliant architecture processes voice packets similarly using either protocol. The software is broken down into two parts, the DSP and microprocessor components. The DSP processes voice data and passes voice packets to the microprocessor with generic voice headers. The microprocessor component is responsible for moving voice packets and adapting the generic voice headers to the NCS protocol. The microprocessor also processes signaling information and converts it from a telephony signaling protocol to IP.

This partitioning of functionality between the DSP and microprocessor provides a clean interface between the generic processing functions -- such as compression, echo cancellation and voice-activity detection -- and the application-specific signaling and protocol processing.

DSP Component (or Voice Processing Module). This software prepares voice samples for transmission over the packet network. Its components perform echo cancellation, voice compression (to conserve cable bandwidth), voice-activity detection, jitter removal, clock synchronization and voice packetization. This unique software, along with TI's programmable DSP technology, provides a comprehensive yet flexible foundation that allows equipment providers to shave months off typical development schedules, resulting in tremendous cost savings and a critical time-to-market boost.

Microprocessor Component. In NCS-based products, the microprocessor component handles detection of various events, reporting of events to call agents, DTMF digit collection and reporting, application of signals and forwarding of audio packets. The microprocessor component in NCS-based products is comprised of the following software modules:

- XGCP Signaling Module (XGCM)
- Digit Collection Module (DCM)
- DSP Interface Module (DIM)

DSP Component

PCM Interface. This interface receives PCM samples from the digital interface and forwards them to the appropriate DSP software modules for processing. It also forwards processed PCM samples received from DSP software modules to the digital interface. It performs continuous re-sampling of output samples to avoid sample slips.

Tone Generator. This generates DTMF tones and call-progress tones under command of the host (e.g. telephone, modem, PBX or telephone switch). It supports U.S. and international tones.

Echo Canceller. This performs G.165 and G.168-compliant echo cancellation on sampled, full-duplex voice signals. It has a programmable range of tail lengths.

Voice Activation Detector. This monitors the received signal for voice activity. When no activity is detected for a specific period of time, the software informs the IP protocol. This prevents the encoder output from being transported across the network when there is silence to save bandwidth. This software also measures the Idle Noise characteristics of the telephony interface. It reports this information to the IP protocol in order to relay this information to the head end for noise generation when no voice is present.

Tone Detector. This detects the reception of DTMF tones and performs voice/fax discrimination. Detected tones are reported to the Host so that the appropriate speech or fax functions are activated.

Voice Codec Software. This software compresses the voice data for transmission over the packet network. It is capable of numerous compression ratios through its modular architecture. A compression ratio of 8:1 is achievable with the G.729 voice codec.

Fax Software. This software performs a Fax Relay function by demodulating PCM data, extracting the relevant information, and packing the fax-line scan data into frames for transmission.

Voice Payout Unit. This buffers voice packets received from the packet network and sends them to the Voice Codec for payout. The following features are supported:

- FIFO buffer that stores voice codewords before payout to remove timing jitter from the incoming packet sequence
- Continuous-phase resampler that removes timing-frequency offset without causing packet slips or loss of data
- Timing-jitter measurement which allows adaptive control of FIFO delay

The voice packetization protocols use a sequence number field in the transmit-packet stream to maintain temporal integrity of voice during payout. Using this approach, the transmitter inserts the contents of a free-running, modulo-16 packet counter into each transmitted packet, allowing the receiver to detect lost packets and to reproduce silent intervals during payout.

Packet Voice Protocol. This encapsulates compressed voice and fax data for end-to-end transmission over a backbone network between two ports.

Control Interface Software. This coordinates the exchange of Monitor and Control information between the DSP and Host via a mailbox mechanism. Information exchanged includes software downline load, configuration data and status reporting.

Real-time Portability Environment. This provides the operating environment for the software residing on the DSP. It also provides synchronization functions, task management, memory management and timer management.

Unsolicited Grant Service (UGS). Cable networks are asymmetric -- the downstream data received is streaming while the data transmitted upstream is either transmitted on a collision time-frame or must get a time slot or "grant." Because requesting a grant can cause significant delay, UGS ensures that cable modems will be contacted at regular intervals without having to make separate requests. The concatenation process mentioned earlier can lighten UGS requirements and increase the efficient use of bandwidth

UGS with Activity Detection (UGS-AD). Upon detection of voice inactivity UGS-AD enables network resources to be diverted to other cable modems and data flows, maximizing the efficiency of all data transmissions.

Microprocessor Component

DSP Interface Module. The DIM is responsible for the interface to the TI DSP-based Telogy Software™ products. The microprocessor communicates with the DSP through a shared memory arrangement whose mechanics are hidden by the DIM. The DIM shields the rest of the microprocessor software from the complexities of the DSP interface.

XGCP Signaling Module. This module is responsible for providing MGCP Embedded Client functionality. It parses and processes each message received from an MGCP call agent. It reports

detected events to the call agent, generates signals requested by the call agent, reports detected DTMF digits, and sets up connections requested by the call agent. This module is also responsible for forwarding audio packets received from the DSP to the packet network interface and forwarding audio packets received from the packet network interface to the DSP.

Digit Collection Module. This module is responsible for processing dialed digits received from the XGCP module. It accumulates all the dialed digits and matches them against the digit map. It reports the results along with the accumulated digits to the XGCP module.

Network Management Module. This module is responsible for providing the management interface to configure and maintain the other modules of TI's Telogy Software™ products. A sample module is provided, but the customer may replace the sample with a custom module. A proprietary Voice Packet MIB is supported because no standard MIB exists.

VII. Fault Analysis

Fundamental to any communications system is the ability to discover, isolate and remedy problems as quickly as possible to minimize or eliminate the degree to which users are impacted. TI's Telogy Software™ products offer a rich set of diagnostic features to accomplish this.

Loop-back Capability. PCM loop-back is used for diagnosing problems related to the telephony side. Pack-send loop-back diagnoses problems related to DSP software performance. Packet-receive loop-back diagnoses problems related to the packet network.

Signal Level Measurement. TI's Telogy Software™ products can perform signal-level measurements on the telephony-side interface for a particular channel. It reports instant and mean values for signal power in both directions.

Packet Network Statistics. The software generates an extensive set of performance statistics for each channel. The statistics include number of transmitted and received packets, minimum and maximum packet inter-arrival times (i.e. jitter measurement), number of invalid packet headers, and number of lost packets. In addition, the Voice Playout Unit reports number of lost voice frames, number of repeated frames, number of idle frames, number of dropped frames, and average packet jitter.

PCM Sample Trace. The DSP software can provide 10 milliseconds of PCM samples (approximately 80 samples).

Memory Trace. The DSP software can provide 40 memory values, starting from the requested location. Memory locations can be from data memory or program memory.

Fax Processing Debug. TI's Telogy Software™ products provide a stream of debug information tracing the performance of fax operations.

Echo Canceller Statistics. The software offers a detailed set of echo-canceller debug statistics.

VIII. Summary

With the merging of telecom carriers, cable operators and Internet service providers, most experts agree that convergence is not merely a trend but inevitable. The potential cost savings, consolidated billing, streamlined network management and overall convenience are too compelling for service providers and consumers to ignore. With buried cable passing hundreds of millions of homes worldwide, it is logical to assume that cable will be front and center as convergence becomes mainstream.

The technical challenges will be overcome as innovation and experience combine to provide cable-based IP telephony solutions that are the equal of the public telephone system. The software architecture outlined in this paper has been field tested in equipment using Texas Instruments' integrated microprocessor and

programmable DSP chipsets. It is an architecture designed to provide equipment manufacturers with a repeatable, “core” starting point to help them develop unique, value-added telephony solutions and bring those solutions to market as quickly and cost-effectively as possible.

About the Authors

Edward Morgan is the Executive Director, R&D of Telogy Networks, A Texas Instruments Company. He is responsible for all voice, fax and data over packet software development. He has 16 years of experience in design and development of digital communications systems at Telogy Networks and AT&T Bell Laboratories. He has a number of patents in the area of digital circuit/packet switching. He received his MSEE and BSEE from Drexel University, and MBA from Purdue University.

Debbie Greenstreet is Senior Product Manager at Telogy Networks, A Texas Instruments Company. She has over 15 years experience in product management, systems engineering and product development with major US and international communications equipment companies. Ms. Greenstreet holds a BSEE from University of Virginia and has performed graduate work in computer engineering.

¹ CableLabs is a research and development consortium of cable operators from across North America and South America. The organization’s guiding concept is interoperability as cable networks evolve from regional to national in scope, and the industry branches into Internet and telephony services.

References

Reference sources for this paper include interviews with Telogy Networks employees and many other internal documents. Additional source materials are:

Barton, Henry. *DOCSIS MCNS vs DVB/DAVIC DVB-RCC: The Case for DOCSIS in Europe*. White paper, May 26, 1999. <http://www.cable-modems.org/standards/DOCSIS/DOCSIS-in-Europe+White+Paper+v1.pdf>

Clarke, Peter. *European consortia back rival cable-modem standards*
EE Times October 28, 1998
<http://www.eet.com/story/OEG19981028S0005>

Hoffpauir, Scott. Broadsoft, Inc.
Presented at Fall '99 VON conference
<http://www.broadsoft.com>

Schulzrinne, Henning. *SIP Overview*
Columbia University
<http://www.cs.columbia.edu/sip>

Cable-Modems.org, online tutorials
What is a Cable Modem?
<http://www.cable-modems.org/tutorial/WhatStandards?>
<http://cable-modems.org/tutorial/13.htm>

Cable Modem Standards and Specifications

Cable Modem Info Center, Cable DataCom News

<http://cabledatcomnews.com/cmhc/cmhc3.html>

ETSI Accepts DVB as European Standard For Data Communication Over HFC Networks. DVB/DAVIC Interoperability Consortium, news release, April 26, 2000. http://www.dvb-davic.org/press_releases/etsi.html

Interim 1.0 Specification

PacketCable; <http://www.packetcable.com>

Overview of DVB-RCCL / DAVIC vs. MCNS / DOCSIS.

DVB/DAVIC Interoperability Consortium, white paper, October 20, 1998. http://www.dvb-davic.org/white_papers/dvb_rcc_wpaper.html

Security in DOCSIS-based Cable Modem Systems; Specs Technology, CableLabs, September 1999
CableLabs. http://www.cablelabs.com/about_cl/SPECS/September_SPECSTECH/tech.pgs/leadstory.html

SpecsNews fom CableLabs, December 1998.

http://www.cablelabs.com/about_cl/SPECS/specnewsdec/news.pgs/story2.html