Developers gain several advantages by buying hardware and software components and using open industry standards.

Open Building Blocks Ease the Creation of Telecom Test Systems

By Gordon Wilkinson

It's vital that mobile and terrestrial telecommunications network providers maintain high standards of quality and reliability; hence, issues like quality of service and system uptime are crucial. In turn, well-designed test and simulation equipment is essential to ensure that a telecommunications network reaches the required QoS level.

Owing to the size and complexity of networks, it's important that test and simulation systems handle many and varied tasks. At this scale, test equipment must be custom-built for the application under test, but the customization need not be as lengthy nor as expensive as it sounds. A wide variety of test and simulation systems can be built from standard building blocks. These systems use open industry standards, which enable developers to buy hardware and software components and speed the completion of their project while reducing the cost and inherent risk.

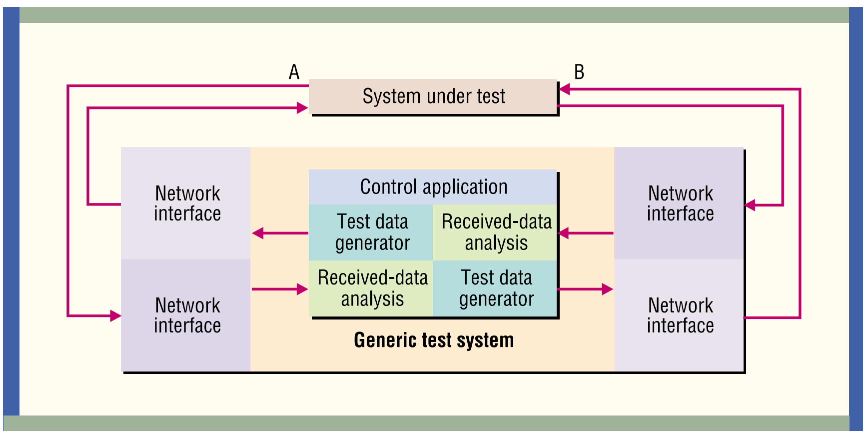

Figure 1. Test data is generated and passed via a network interface to the system under test. This data is operated upon by the system under test, sent to the test system, and then compared with the data that was originally sent.

Because QoS is so heavily influenced by the behavior of the system while it's under load, it's essential that the test system be able to create a significant load and measure its individual effects. Such load tests allow system developers to analyze and adjust their queuing and prioritization algorithms, as well as jitter buffer management, under controlled conditions, depending on how the test system has been configured to apply the load.

THE GENERIC TEST SYSTEM

Figure 1 shows a generic telecom system connected to a simple two-port test system, although the same

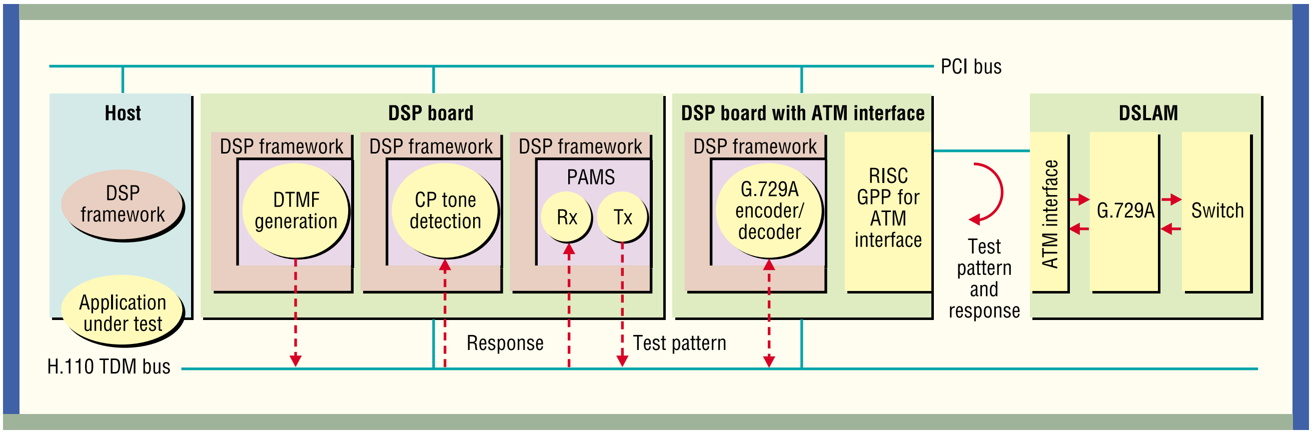

Figure 2. The DSL access module (DSLAM) test system shown uses a DSP board performing tone detection and generation in order to set up and tear down calls to the module and the Perceptual Analysis Measurement System (PAMS) to generate test data and compare the resultant data. A second DSP board with an ATM interface is used to connect to the DSLAM.

principles apply to multiple-port systems. Connection to the system under test is made via network interfaces, with the types of interface depending on the system. For example, a voice-over-Internet Protocol gateway is likely to have side A interfacing with a PSTN using T1 or E1 lines and side B interfacing with the Internet or an intranet with an IP connection.

The test system works by generating test data and applying it to either port. The test data can be simulated calls that exercise the services available using a variety of signaling operations. Alternatively, it can be tones or simulated speech along the voice channel itself. In either case, the resultant output from the system under test is received by the test system and analyzed for the correct response to the applied test data.

The main tasks of the test system are to simulate the operation of a real network at both the inputs and the outputs of the system under test and to verify that the stimulus provided was handled correctly. The overall operation of the test system is managed by a control application that schedules the tests and receives the results of the analysis. This separation of control functionality and the functions within the speech channel, like generating test signals, results in a natural division of tasks between processors (see “Teaming DSPs with GPPs,” page 25).

The main tasks are to simulate the operation of a real network at the inputs and outputs of the system under test and to verify that the stimulus provided was handled correctly.

AUTOMATED TESTING OF SPEECH QUALITY

Early test systems were able to use simple tones at a range of frequencies to verify the frequency response of the system under test. Modern systems require some additional sophistication. For example, testing a mobile-telephone base transceiver station (BTS) or VoIP gateway calls for the consideration of speech compression algorithms, such as G.729A and G.723.1 for VoIP or Enhanced Full Rate (EFR) for GSM speech, which may also be required to form part of the test system, depending where in the network the tests are to be made.

These algorithms are based on models of the human voice, so a tone isn’t a suitable test signal. Instead, a test pattern of real or simulated speech is used. These test patterns are typically standard simulations stored in a WAV file and passed to the DSP for preprocessing (speech coding, if required) and storage and for comparison with the received signal (which may require speech decoding).

Various algorithms are available for automated testing of speech quality. They're designed to replace the mean opinion score (MOS) tests, which required real human beings to listen to the quality of the received speech and award a score from 0 to 5, with 0 the worst quality and 5 the best. Clearly, this approach isn't practical for load-testing a network, although MOS is still often quoted as a figure of merit for a codec algorithm and used as a guideline for its eventual performance in a system.

The most commonly quoted algorithm for automated speech quality analysis is perceptual speech quality measurement (PSQM). Here, the artificial speech must adhere to ITU-T P.50 or, if real speech is used, to ITU-T P.830. These standards ensure a long-term and short-term average spectrum, an instantaneous amplitude distribution, and a syllable envelope.

In short, the PSQM algorithm subtracts the results of a perceptual model of the test pattern (the way the ear would receive the signal) from the results of a perceptual model of the received signal. The output of the subtraction is assessed

The assessment amounts to a series of mathematical operations.

using a cognitive model that simulates the way the brain would understand the speech. The assessment amounts to a series of mathematical operations on the sampled speech and is ideally suited to digital signal processing. The values obtained from a PSQM test are 0 (perfect quality) to 6.5 (worst quality).

One drawback of PSQM is that it doesn't take into account delays in transmission, which are a major issue with VoIP. The packets of speech transmitted using VoIP often take different routes and therefore can arrive out of order at their destination. This problem is handled using a jitter buffer, a section of memory corresponding to a given length of time. The buffer orders the packets correctly before they're played. The longer the delays in the network, the larger the necessary buffer must be to provide acceptable speech quality. The delay tends to become an issue to the listener above about 200 ms, and it's therefore useful to include this factor in the overall test measurement.

PSQM+ was developed by Royal KPN N.V., the largest telecommunications company in the Netherlands. The Perceptual Analysis Measurement System (PAMS), developed by British Telecom, also takes delays into account.

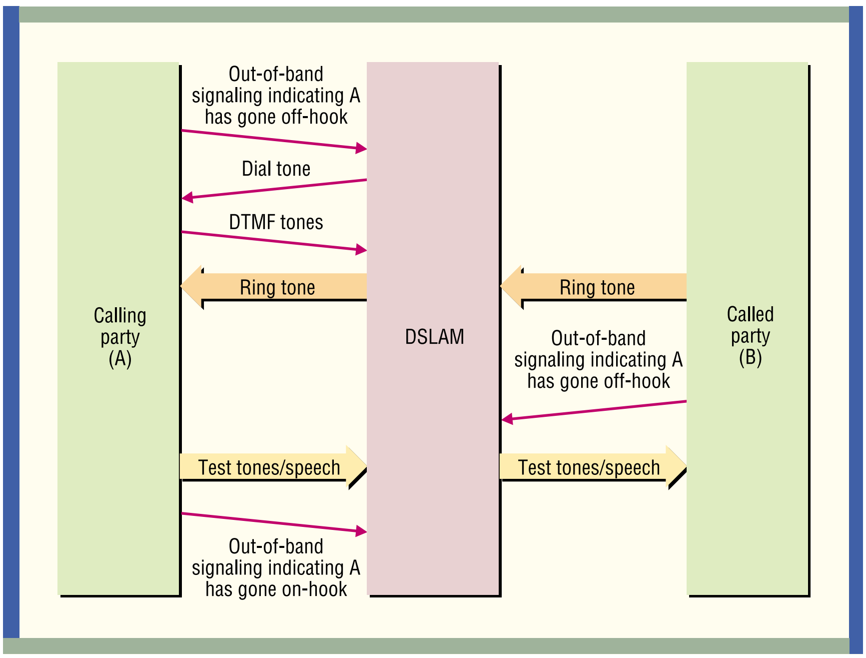

Figure 3. Calls are set up and torn down in much the same way as in a real network. To do that, the test system must simulate a calling and a called party.

OPEN STANDARDS

System development, all the way down to writing DSP algorithms and managing many instances of them, is a mammoth task. To make life easier, Texas Instruments has introduced the TMS320 DSP Algorithm Standard, which allows DSP algorithms to be integrated easily into a framework capable of executing and controlling many algorithms simultaneously in real time. The standard allows a system developer to buy compliant algorithms from a developer—a major benefit, particularly for optimized algorithms on complex devices such as today's very long instruction word (VLIW) DSP architectures.

TI provides a function-naming convention, including names for some mandatory functions. The names are located using a virtual table so that a DSP framework can

Text continues on page 26

dynamically allocate them. The DSP framework provides the means by which the algorithms are controlled and scheduled and presents a structure by which the algorithms are called. It also offers communication paths among the serial ports and TDM and between the DSP and the host for passing data and control messages.

A corresponding framework is needed on the host as well, providing an API to control the operations on the DSPs via the control path to the DSP framework. This framework has the benefit of abstracting the programmer from the operations of the DSPs. DSP and host frameworks are available commercially.

Figure 2 shows a DSL access module (DSLAM) test system utilizing an ATM interface board with DSP capability and an H.110 connection to a DSP board. The DSP board is used to run several different algorithms: call progress tone detection and generation; in-band call control signaling; and a PAMS algorithm for test signal transmission, receiving, and comparison. The DSP on the ATM board handles speech coding and connects to the system under test through the ATM interface.

The host framework abstracts the programmer from the DSPs' operations.

A general-purpose RISC processor on the ATM board manages the ATM interface, receiving packets from the on-board DSP handling speech compression and turning them into ATM cells for transmission to the DSLAM. The DSLAM is configured to loop back the call at the switch, meaning that the speech is decoded and encoded before being reconverted to ATM cells and returned to the test system. On the return path, the GPP deconstructs the ATM cells and passes them back to the on-board DSP for decoding.

Overall control is from the host, which runs a test application—proprietary protocols that communicate with the GPP to give it configuration information and APIs that tell the DSPs which operations to perform and which time slots on the H.110 bus are to be used for input and output. These APIs are part of the host framework discussed above. The operation is such that the test application simulates a calling party, A, and a called party, B, with a sequence of events similar to that of a conventional call, as shown in Figure 3.

The first step is to send out-of-band signaling information to indicate that A has effectively gone off-hook and is ready to make a call. The DSLAM responds, indicating that a path is clear by generating a dial tone, which is inserted into the payload of a cell and sent back to A.

The GPP on the ATM board in the test system then extracts the payload and sends the dial tone to the on-board DSP. In this case, since the payload isn't speech data, the DSP operates transparently and doesn't decode the received data. Instead, it places the data directly onto the H.110 bus, in a time slot that is being monitored by the call progress (CP) detection algorithms running on the DSP board.

The host is sent a packet by the DSP performing the CP detection to indicate that a dial tone has been detected and passes a string of digits to the relevant DSP, indicating which DTMF tones to generate. The digits are passed across H.110, via the DSP on the ATM board and across the ATM interface. The DSLAM then makes a connection to B and generates a ring tone to both A and B.

Once the CP tone detector recognizes the ring tone on B's incoming H.110 time slot, it indicates this to the host. The host now signals that B has gone off-hook through out-of-band signaling via the ATM GPP.

When the CP tone detector for A detects that the ring tone has stopped and that B has therefore answered the call, it indicates this to the host. The host then passes the test file to the PAMS DSP, which sends it out in A's outgoing H.110 time slot, this time for G.729A encoding on the DSP ATM board before transmission to the DSLAM.

The signal stream is looped back in the DSLAM and returned to the test system for G.729A decoding and placement in B's receiving H.110 time slot, where it's picked up by the PAMS DSP for comparison with the transmitted signal. The results of the comparison are sent back to the host when the file is completed and the call is closed down.

Measurements of interest can be separated into two types: signaling and speech quality. Signaling measurements include how long the DSLAM takes to answer the call and connect the two parties, that is, how quickly the DSLAM takes to detect or generate individual tones. To some extent, speech quality measurements are a function of how the DSLAM handles its ATM cells.

Note that the test system typically sends the DSLAM tens or hundreds of simulated calls at once in order to assess its QoS under load.

Ideally, the calls stress the prioritization and queuing algorithms to the point at which the DSLAM is forced to throw away packets. (Discarding packets can occur under intense loading.) The tests also measure the performance of the codecs in the DSLAM, including the jitter buffering and the switching circuitry. ◆

Gordon Wilkinson (gordon.wilkinson@motorola.com) is a DSP systems architect with the Motorola Computer Group in Loughborough, England, following Motorola Inc.'s acquisition of Blue Wave Systems. His main focus is wireless and gateway application-enabled platforms. He has 12 years' experience in the embedded computing industry.