Testing Digital Video

The 1997 Digital Video Test Symposium

Author: Dragos Ruiu

Discover the latest techniques and products for real-world video testing

Testing Digital Video

The video industry is making a monumental transition from analog broadcasting to digital data delivery, and the effects will be wide reaching socially as well as technically. Led by satellite service operators—with cable operators, terrestrial broadcasters, and wireline and wireless telco carriers following hot on their heels—there is a stampede to be first to market with digital service bouquets. But the path to this new technology is full of pitfalls. In this paper we will examine techniques for avoiding some of those pitfalls. We'll look at the measurements and tests that can be used to develop, install, maintain, and troubleshoot these new digital video (DV) systems.

We begin by examining the kinds of impairments and problems that can occur on DV networks. First we look at the generic issues that affect the compression and transport system common to all these systems: ISO MPEG-2. Then we consider system-specific issues related to some of the flavors of video transmission systems currently being deployed around the world. Finally we examine some of the emerging techniques for testing the MPEG encoding and decoding conversion from analog to digital and back.

A caveat applies, however, because in many ways these systems are different from any we have dealt with in the past--and as a result, we are just now learning what it means to test the new applications and devices. The road to building reliable high quality compressed digital video networks will undoubtedly be full of surprises. But, with careful planning, it should be possible for an operator to bring up reliable digital services from day one—a crucial goal, because these new digital services must compete with the trusted, well quantified and well understood analog equipment and services. Digital technology opens up a realm of possibilities for interactive services and applications, but in the early days, digital video equipment purchases will be justified by the same source of revenue as the analog systems: television entertainment. Therefore, it is mandatory that the DV services offer the same, if not better, levels of service and quality that subscribers have come to expect from the analog TV that we have refined over the last fifty years.

Network Architecture

The systems being deployed for DV delivery and transport fall into several categories, each with some unique implementation challenges and testing issues:

- Satellite systems: Direct broadcast satellites have been early to market in delivering a high quality and large quantity of digital programs to most of the geographies of the world. Two major systems currently exist for satellite data delivery, the DSS system designed by Hughes and a consortium of other companies, and the DVB-S system developed by the European Union's Digital Video Broadcasting Project. Both systems add the infrastructure to the ISO's MPEG-2 broadcast video compression system to realize full one-way television broadcasting from satellites to small consumer-owned dishes and digital decoders. Service-specific variants of the DVB-S system exist, and they differ primarily in the delivery of EPG, conditional access, and system data.
- Cable systems: Most, if not all, terrestrial cable operators are looking at the new compressed DV technology to give their hybrid fiber-coax architecture a capacity and service boost. Four basic transmission systems for this kind of network have been standardized by the ITU, of which three are in common use. The DVB Project has produced two variants of a digital cable transmission system, called DVB-C, which can be used to incrementally expand current cable systems on a channel-by-channel basis. One variant (ITU-R J.83 Annex C) for 6 MHz NTSC channel slots is being deployed in Japan and the rest of Asia, while another (ITU-R J.83 Annex A) for 8 MHz PAL channel slots is being deployed in Europe and other PAL-based nations (such as some South American countries). Both use quadrature amplitude modulation (QAM). Another QAM-modulated variant (ITU-R J.83 Annex B), based on specifications from the two dominant North American cable equipment manufacturers, General Instruments and Scientific Atlanta, is being deployed in North America.

The North American specification differs primarily in the error correction scheme used and the format of the electronic program guide information. A fourth variant (Annex D), based on the vestigial side-band (VSB) modulation used in North American terrestrial digital transmissions and proposed primarily by Zenith, has not seen much commercial adoption. Vendor-specific sub-variants of these systems also exist.

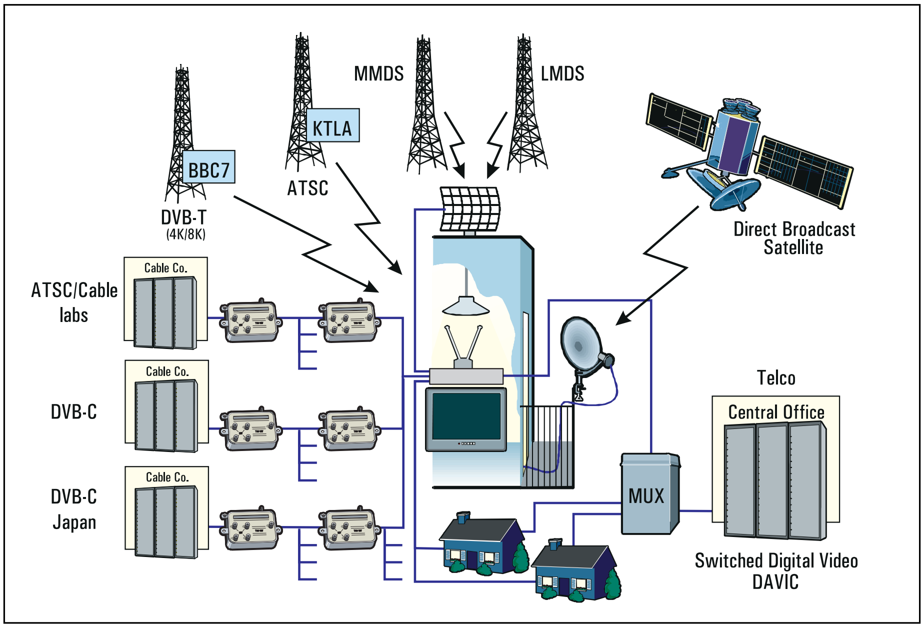

Diagram 1:
DV Systems

- Terrestrial Broadcasting: The DVB project has specified a coded orthogonal frequency division multiplexing modulated (say that ten times fast!) or COFDM-based terrestrial RF broadcasting transmission system. This system, DVB-T, is being deployed in Europe. It has two variants, the so-called 4K and 8K variants. The 4K variant is being deployed mainly in the UK and is upwardly compatible with the 8K variant being deployed in the rest of Europe and parts of South America and Africa. In North America, the FCC's Advanced Television Standardization Committee (ATSC) has specified a VSB-based terrestrial transmission system, and local broadcasters have been assigned additional spectrum and begun limited trial broadcasts using this system in both normal NTSC resolution and HDTV. Two other, line-of-sight, microwave terrestrial transmission systems are the cellular-like Multipoint-Multichannel Distribution System (MMDS) and Local Multipoint Distribution

System (LMDS). These systems are being used by new entrants into the video services arena (and for data applications), and they have been most publicized on the East Coast of the U.S. and in the Los Angeles area, where Pacific Bell has rolled out an MMDS-based TV service.

- Switched Digital Video: The star architecture SDV is being examined and extensively trialed primarily by telephone companies looking to leverage their extensive twisted-pair copper infrastructure. The Digital Audio Video Council (DAVIC) has developed some SDV standards, but vendor-specific variants are still the norm.

DV Systems

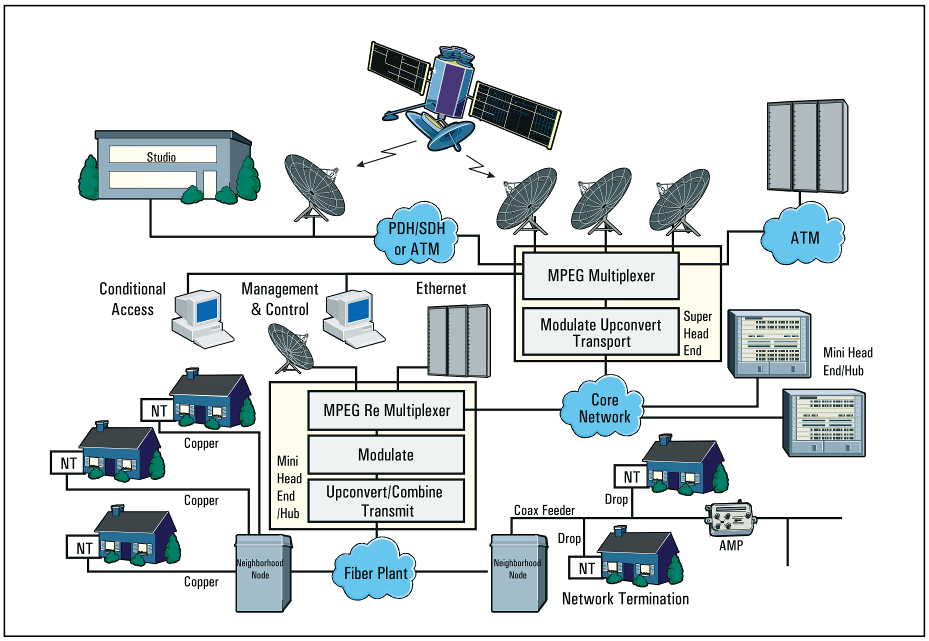

Diagram 2: Generic Network Architecture

While each of these systems differs radically in the access network architecture used for final delivery of information to the consumer, and in some of the formats used for control and program guide information, the core architecture and the equipment used to deliver the services share a common architecture in the interior of the network. We will now examine the architecture with an eye towards identifying good test and measurement points.

Programming content is sometimes delivered using existing analog contribution networks and equipment (satellite, terrestrial RF, and land-line) and compressed at the service provider's facility. Increasingly, though, the signals are being compressed at the source and transmitted digitally throughout the entire path, because the compression results in transmission cost savings. With the larger bandwidths involved with video, ATM transport and services are commonly used, taking advantage of ATM switches' high bandwidth switching capacity. For interactive content and commercial insertion, video servers are often linked to the headends via ATM. Cable headends have limited physical space, so often the server farm is in a different physical location (sometimes even in a different city!) and ATM switches are used to route the program streams to the headend.

In typical configurations, the servers deliver single program transport streams (SPTS) without program system information (PSI) tables and electronic program guide (EPG) information. The multiple SPTSs are combined and multiplexed by a multiplexer at the master head end into a multi-program transport stream (MPTS) with PSI and EPG content. Typically, it is at this point where EPG and conditional access (scrambling keys for NVOD, VOD, and PPV) are integrated into the transmissions from workstations and servers connected to the multiplexer via LAN networking. This combined stream is fed as 188 byte transport packets (or sometimes null padded 204 byte packets) over a short haul parallel or serial interface to a modulator. While various proprietary interfaces exist, the common configuration seems to be a DVB-specified synchronous parallel interface (DVB-SPI) or an asynchronous serial interface(DVB-ASI). The SPI interface seems to be preferred for short cable runs, and the coaxially cabled ASI interface is used for longer physical distances.

Usually the encoders and multiplexers are configured in a redundant configuration, with one or more hot-spare encoders and a fully redundant multiplexer, so that equipment failures will cause minimal service outages. In satellite and cable systems, the modulators take the input transport stream and add forward error correction information to compensate for transmission errors. The resultant RF streams are then upconverted to the appropriate frequency slot, combined, and sent to the appropriate laser or transmitter. In some systems SONET or SDH is used as the core transport network, and the RF modulation does not occur until later.

In most network configurations, this master signal feed is distributed via existing telecommunications infrastructure to regional distribution nodes or transmitters. At these nodes the signal is redistributed. In more complicated network scenarios, these regional nodes will re-multiplex the signal to add local content (or potentially do local ad and content insertion), and a further multiplexing step will occur before the final modulation. In HFC and SDV architectures, a neighborhood optical network unit (ONU) node converts the optical transmissions into electrical ones. In SDV systems more multiplexing and channel selection occurs at the neighborhood nodes.

To test this system, we should consider each of the transmission systems and devices as a point of failure. An appropriate test strategy would be to verify the operations of each of these devices in turn. Testing of cabling and RF measurements can and should be made throughout the system. MPEG-level measurements need to be made up to the last point at which the digital information is modified or multiplexed.

It is appropriate to make MPEG transport stream measurements wherever multiplexing of the stream occurs. Because multiplexing does not affect the video and audio contents, it does not make much sense to make measurements of the video and audio information in the subsequent network steps after the encoder or server. Similarly, after the modulators, bit and transmission errors will have equal probability of affecting all the bits, so transport-level tests should give way to RF and spectrum measurements. Let’s examine each of these domains separately, focusing initially on the MPEG transport system.

Transport Measurements

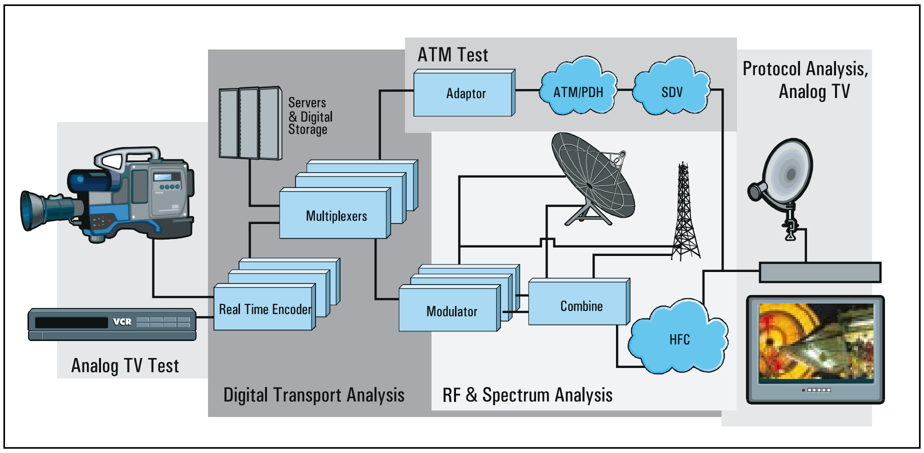

Diagram 3: Measurement Domains

There are a number of tests that should be made at the transport level. The transport stream itself has some control structures that should be examined to verify the health of the system. The MPEG 2 system specifications (ISO-13818 section 1) defines a flexible multiplexing system that can mix multiple channels together.
The information is transmitted in fixed-length, 188-byte packets. These packets consist of a 4-byte header, containing control information, and a 184-byte payload in which the video, audio, and data information is transmitted. The header information provides some valuable information for testing.

Diagram 4: Transport Stream Packet Format

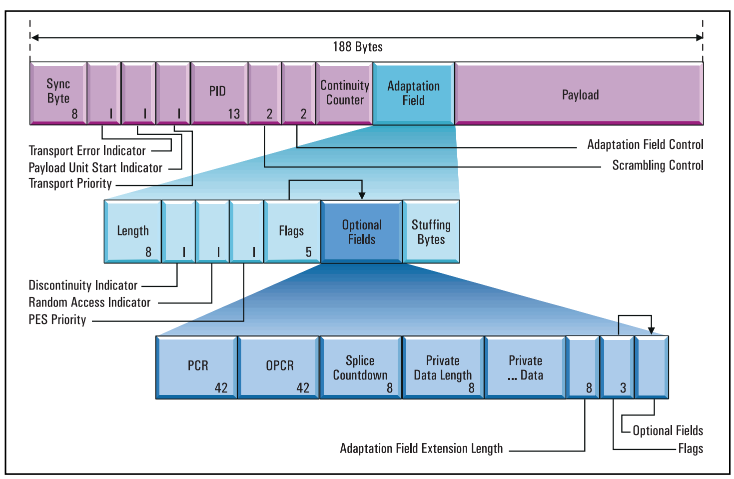

Each packet begins with a fixed value (0x47) that can be used for synchronization and framing. If packets are being received with improper sync bytes, we know that something is seriously wrong. The transport stream consists of a steady stream of these packets. When the system does not have data to transmit, null packets are transmitted to fill the space. The individual channels and streams are identified by a packet identifier (PID) field in the header. Each video and audio stream is assigned a fixed PID number for the duration of the program. This PID value can change as it goes through a multiplexer. The PID streams form virtual channels which are multiplexed together and transmitted inside the physical link.

There is a lot of information in the transport stream header that can be used for testing purposes. The first is the transport error indicator (TEI). This is a bit reserved for the transfer of management information throughout the MPEG transport. When a piece of MPEG equipment, particularly a multiplexer or re-multiplexer, detects errors on incoming signals, it is supposed to set the TEI on traffic that it passes to subsequent equipment in the transmission chain. Incoming errors can be detected in a number of ways, ranging from error correction code detection to identifying format errors. For troubleshooting purposes, examining the streams to see if there are any TEI bits on in the traffic can be a useful technique.

Another field that is very important for testing purposes is the continuity counter. This field increments modulo-16 for every packet transmitted on a particular PID. Test equipment can detect lost or dropped packets by looking for jumps in this counter.

Transport Error Indicator (TEI)

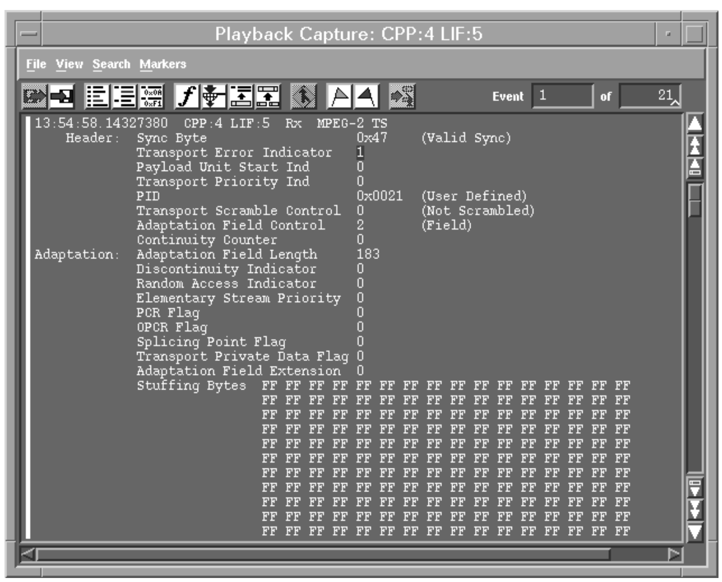

Diagram 5:
PCR Jitter
& Interval

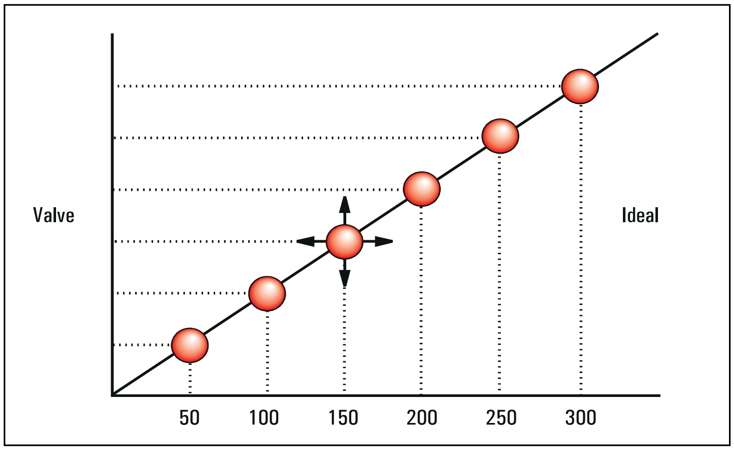

PCR Jitter

The transport stream header also contains the adaptation field, an optional extension header that can conditionally be included into the PID stream. The adaptation field contains a very important part of the MPEG transport information, the timing information for synchronizing the receiver to the transmitter. So that frame rates match the display rate to the transmission rate, MPEG encodes a 27 MHz master clock into each stream. This clock is called the program clock reference (PCR). For synchronization, a sample of this numeric clock is taken at the encoder and inserted periodically into adaptation fields on the output stream. The ISO specifies that this occur at least once every 100 milliseconds. The DVB project has tightened the specification to 40 milliseconds to ensure that clock drift doesn't affect the generation of the color sub-carrier in the decoder.

Periodically, then, the value of this clock can be found in TS adaptation fields for each program. All the programs run asynchronously to each other to avoid the need for time-base correction and global synchronization. The clock is fundamental to the operations of the decoder—all the other timestamps such as presentation time-stamp (PTS) and display time-stamp (DTS) are stated in terms of this master PCR clock, and most of the decoding timing is derived from this clock. Because of the importance of this clock, its stability is an important test concern. If too much jitter and wander exist in the clock, then the decoder will experience buffer overflows or under-runs, owing to its limited amount of memory (the decoder in MPEG is optimized to reduce the number of expensive RAM components in the set-top).

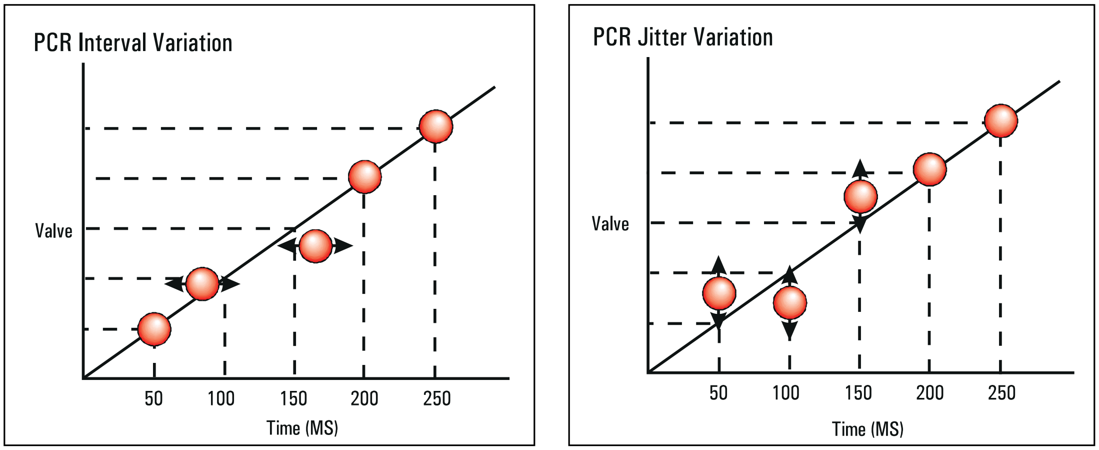

Diagram 6: PCR Interval/ Jitter Variation

In the case of excessive PCR errors, the decoder will exhibit audio and picture breakups.

In measurement equipment, the challenge of measuring jitter is to determine what the "ideal" time base is from which the samples are deviating. This is typically done by collecting a larger number of samples and approximating the ideal timing line from these samples. The "ideal" time line is approximated by using a best-fit algorithm on the data points. In the HP implementation, this best-fit line is calculated on a continuous basis from a sliding window of 100 samples. Jitter is then measured as the deviation from this continuously calculated best-fit line. HP has a number of patents pending relating to these techniques for PCR jitter measurement.

Typically the jitter tolerance of the decoder is a fixed value determined by its design. This jitter sensitivity gives network designers an overall jitter budget for the entire network—all of the jitter introduced by each successive network hop or piece of equipment cannot total more than the jitter allowable at the decoder end-point. To determine the sensitivity of the decoder, streams with progressively more jitter can be transmitted until the decoder starts to fail. This overall jitter performance is an important step in basic network design. While not all the jitter may be additive (some jitter introduced from multiple elements can cancel each other out), for safe design purposes it is best to assume the worst-case scenario and treat all jitter as the sum of the individual jitter measured through each device hop in the network.

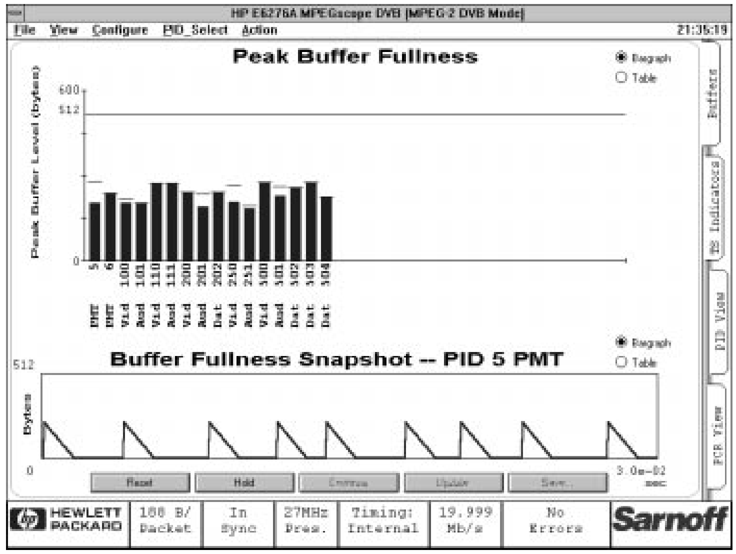

Buffer Level Measurement on MPEGscope DVB

Errors in the PCR can come from two areas, instability in the original clock, which manifests itself as incorrect values of the PCR sample (see Diagram 6), and jitter introduced by the transport and multiplexing, which manifests itself as PCR samples arriving at the wrong time. This second error can occur from packets encountering variable delays in transmission systems such as ATM or from things such as buffer contention between multiple streams in servers and multiplexers.

The PCR is not necessarily continuous. Many network implementations today store video streams on servers with the PCR values pre-calculated as opposed to calculated on the fly in real time. This implies that when switches are made from one program to another (even potentially in commercial breaks), the PCR time base also switches. For this situation, a 1-bit indicator in the adaptation field, the discontinuity indicator (DI), is used to indicate time-base switches. PCR jitter measurements in instrumentation lose meaning across discontinuities (see the discussion of jitter calculation above). For the algorithms used in HP instrumentation, the time base can be reacquired accurately enough to make measurements with some confidence after receiving about 10 samples. This problem is simpler in the decoder, which typically needs only to reinitialize a counter to get the new time base.

In addition to introducing jitter in the PCR, clumping and spreading of packets caused by contention can affect another part of the decoding process, the buffer modeling.

MPEGscope
HP E6277A
PSI Table
Decoding

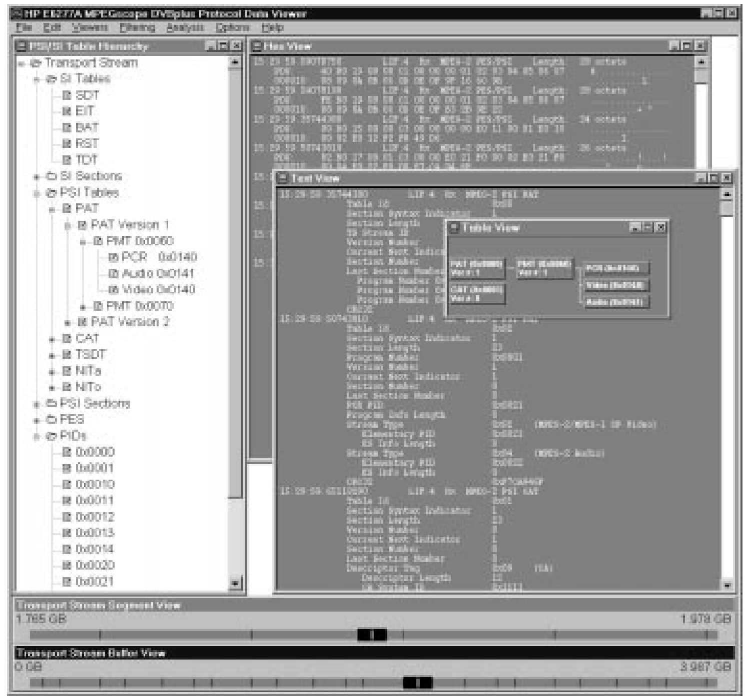

The stringent memory cost requirements of the set-top affect the amount of data that can be buffered. We said earlier that the set-top can only absorb and smooth a limited amount of jitter—because jitter elimination requires buffering, and the more jitter, the more buffer memory needed. The MPEG standard defines a memory utilization model—the so called buffer model—that limits the amount of RAM needed by the decoder. This buffer model places limits on the sizes of the data bursts that can be sent to the decoder. There are three separate buffer models, one for transport packet buffering, nd one each for video and audio data packets.

The transport buffer model limits the number of packets that can be transmitted in a fixed period of time on the same PID. Transport buffers are affected by the multiplexing components in the network. The audio and video buffers limit the size of the audio and video packetized elementary stream (PES) packets that can be transmitted. The audio and video buffer models are affected by the encoding process only, not by the transmission or multiplexing component, while the transport buffer can be affected by the spreading and clumping that can occur as the different programs' packets are interleaved together in a multiplex.

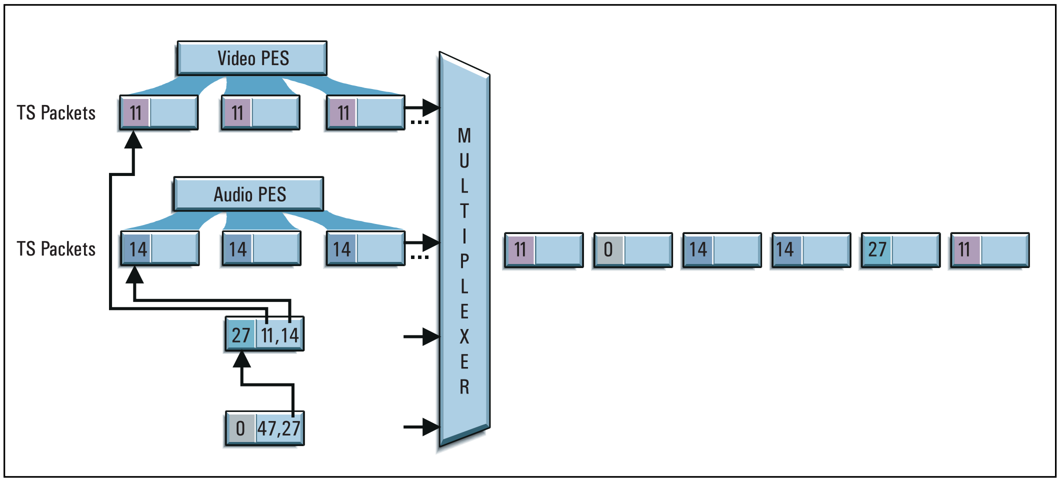

Diagram 7: MPEG Multiplexing

PSI Tables

To identify and separate the different programs being transmitted on separate PID streams, special control data structures are transmitted on reserved PIDs. The root "anchor" of these data structures is the special reserved PID of 0. This PID is never used for programs; it always contains periodic transmissions of a special data structure called the program association table or PAT. The PAT is the initial point at which the set-top can decipher the incoming stream. In the PAT is transmitted a list of other "special" reserved PID numbers. These special PIDs contain periodic continuous low bandwidth transmissions of special data tables called program map tables or PMTs. There is one PMT transmitted for each program or channel being broadcast. The PMTs contain the list of PIDs for the audio, video, data, and PCR streams of each program.

A similar system is used to transmit scrambling keys. Like PID 0 for the PAT, PID 1 is reserved for the conditional access table, the CAT. The CAT is transmitted continuously at a low bandwidth and lists PIDs that are used to transmit entitlement management messages (EMMs). EMM is another name for scrambling keys, which typically identify the addressable decoders that are allowed to receive a particular "enhanced" or non-basic service for which the consumer has paid. The set-top receives the list of addresses, and if it finds its own address in the list, it will enable viewing of that service for the consumer.

These tables structures are specified in the ISO specification and are known collectively as the program system information (PSI) tables. Testing these tables and data structures is important to identifying the basic health of MPEG transmissions, because corruption in these tables will lead to serious decoding problems. The tables themselves are protected against receiving corrupted information by CRC checks, which can be used to discard corrupted information. Since the tables are sent periodically and repeatedly, losing one table is not a serious error; the information will be transmitted again. However, it is important to ensure that equipment designs do not send out table updates either too frequently or too sporadically. The PSI bandwidth consumed typically is less than 100 kbps on links carrying in excess of 20 Mbps.

Another aspect of PSI table testing relates to updates in the tables caused by additions and deletions to the program lineup or other program changes. When changed tables are received, the system increments version numbers in the tables. Using test equipment to look for appropriate version number changes can be a useful troubleshooting technique for verifying changes in the programs being carried on an MPEG multiplex.

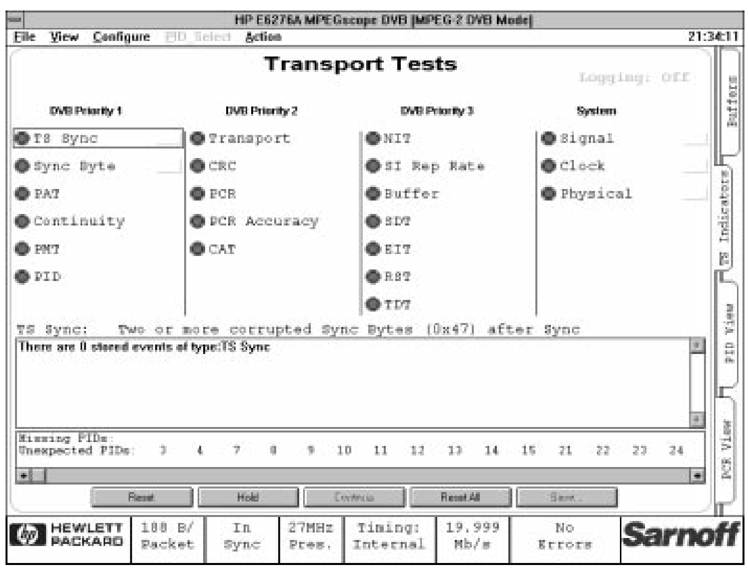

Error Screen Showing Missing PIDs

Diagram 8: Data Tables Transmitted by MPEG DVB & ATSC

ISO Table	PID #	Description
Program Association Table	0x0000	Associates Program Number & Program Map PID
Program Map Table	By PAT	Specifies PID Values for Programs
Network Information Table	By PAT	Physical Network Parameters
Conditional Access Table	0x0001	Associates EMMs streams with PID

DVB-SI Table	PID #	Description
Network Information Table	0x0010	Physical network info for this & other delivery systems
Bouquet Association Table	0x0011	Names of & lists of service bundles
Service Description Table	0x0011	Service data, names, service provider...
Event Information Table	0x0012	Program info on current future prgrms on this & other netwks.
Running Status Table	0x0013	Used for rapid event status updates
Time & Date Table	0x0014	Present time & date
Time Offsett Table	0x0014	Shift in time-zone for current location
Stuffing Table	0x0010-0x0014	Invalidate existing tables at network boundaries

ATSC Table	PID #	Description
Master Guide Table	0x1FFD	Data stream contents, time, PID list, channel grouping
Additional Guide Data Table	0x1FFD	Transmission channels, Program Guide Map, Default
Override Records		
Special Program Guide	by MGT	Additional program guides
Channel Information Table	by MGT	Non changing info, PID #'s, physical channel...
Event Information Table	by MGT	Program info (event title, start time...)
Descriptive Information Parcel	by MGT	Detailed descriptions of Channels & Programs
Private Information Parcel	by MGT	Private/Proprietary description info
Carrier Definition Table	PAT Network PID	Carrier Frequencies
Modulation Mode Table	PAT Network PID	Specifies modulation systems for each transmission syst.
Satellite Information Table	PAT Network PID	Positional information and transponder count
Transponder Data Table	PAT Network PID	Polarizational and waveform description
Network Text Message	PAT Network PID	Text names of transmission systems, currency, rating syst.
Program Identifier Table	By PMT	Creator information about program contents

The PSI tables are used by test equipment to make measurement setup easier for users. Rather than forcing users to look up and manually enter the PID numbers on which they wish to make measurements, the test instrumentation will give users a pre-configured list of PID numbers from the PAT and PMT table configuration that it has detected on the link. This list of valid PID numbers also can be used to look for improperly configured tables or multiplexers, because the test instrumentation can automatically notify the users of any "illegal" PID numbers detected in the traffic but not referenced to in the PMT tables.

Additional tables besides the PAT and PMT tables are transmitted to convey information about the system itself and to deliver more information about the programs. (See chart.) Most of this table information is protected against delivery of corrupt information by a cyclical redundancy checksum (CRC) field that can be used by the decoder to identify and discard information corrupted in transport. Test instrumentation can be used to look for CRC errors to identify link failures or misbehaving equipment.

Data corruption in the control tables of the MPEG link can be devastating, so the CRC checks are used to protect this vital information. Calculating CRC doesn't impose a large performance penalty because the volume of table data is relatively small compared to the volume of audio and video stream data. Unfortunately, no similar protection exists for these voluminous audio and video streams.

Bit Error Rate

The compressed data structure of MPEG is very dense. The process of compression, which removes a lot of the redundancy in the information stream, also leads to increased sensitivity to errors. Each bit in the compressed information stream is more valuable than a bit in the uncompressed stream. Unlike transmission errors in analog TV, each MPEG transmission error can corrupt a portion of the screen for several frames. It is even worse than that, because to achieve high information density, the MPEG audio and video streams are encoded using variable length codes (VLCs). VLCs are a very efficient format, because short codes can be used for information that occurs often, saving bandwidth. But using VLCs means that if one bit is corrupted, the entire bit stream is lost until the next re-synchronization code is received. The decoder identifies dynamically where the current code stops and the next one begins as it decodes the VLCs. Whether one bit is changed or 100 bits are changed, the effect is the same—the decoder loses the input data stream until the next synchronization point is received. The synchronization points in MPEG are the packet start codes, and the smallest unit of data transmitted with a start code is a slice, or sequence of adjacent 8x8 pixel blocks.

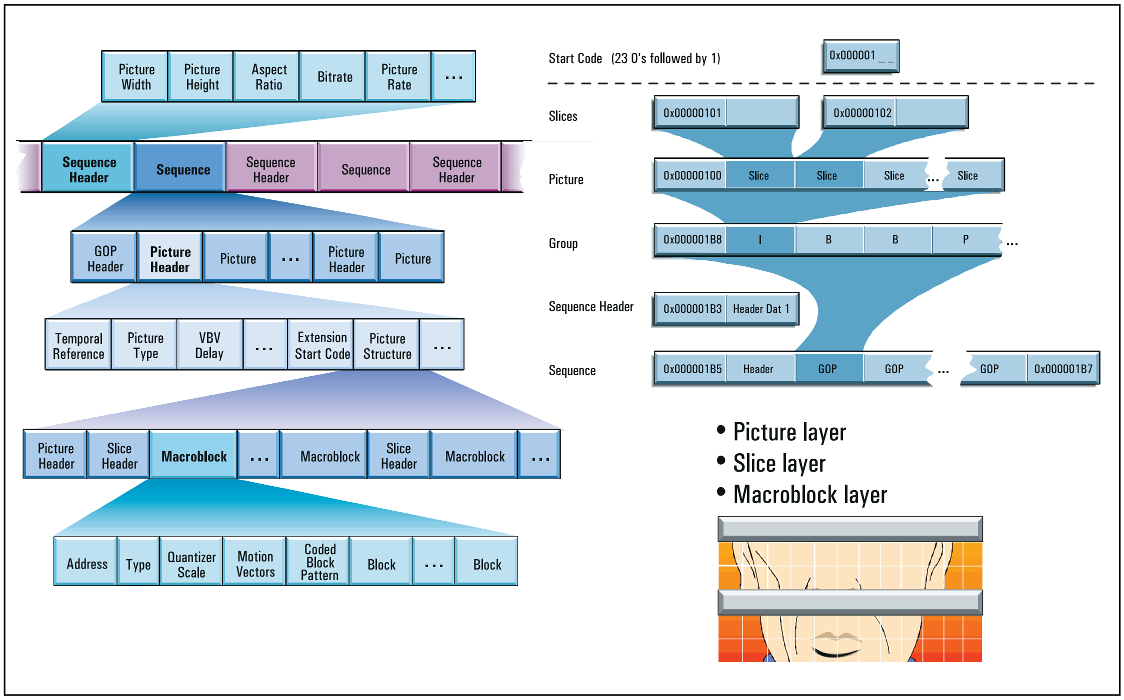

Diagram 9: MPEG Bit Structure & Start Codes

The error sensitivity of MPEG is unusual thing. MPEG doesn't really care how many errors occur in a row, because any error leads to a whole piece of a picture or audio frame (and potentially multiple frames due to inter-frame compression) being lost. Because of the system's extreme sensitivity to errors, measuring the bit error rate (called a BER Test, or BERT) of the MPEG transmission links is very important. It is also important to measure the periodicity of the errors, because although MPEG doesn't care if one bit or 100 are corrupted, we do care how often these error bursts occur. If they occur too often, the picture and audio glitches become annoying to the consumer. One of the important tests in commissioning MPEG links is to transmit a pseudo-random binary sequence (PRBS) on the link while it is out of service, and to measure how often deviations from this pattern occur, in order to determine the BER of the link. The DVB Project Measurement Group has specified a standard PRBS pattern (2e23) to use for these kinds of BER tests.

MPEG systems transmit lots of forward error correction (FEC) information along with the transport packets to protect the fragile MPEG data against corruption. The objective of the FEC system is to improve the error performance of a lossy medium such as satellite QPSK transmission to the point at which it looks like an error free medium (error rates better than 10e-10, or about one glitch per program per day). FEC information is a valuable source of test information, because looking at how hard the decoder is working to compensate for transmission errors can give valuable information about the number of errors on and the health of a link. Extracting FEC information about the number of errors detected and corrected is a handy way to do in-service BER tests without disconnecting the multiplexer and injecting a PRBS pattern. Some vendors include a special test mode in their IRD/set-top to enable installers to get the BER information from the receiver module. Another method of doing in-service BER tests is to transmit a PRBS sequence in the payload of the null packets transmitted to pad the transmissions to the final bit rate.

Checking the error rate of the links is sufficient to test the delivery of the packets from the multiplexer to the decoder, but the link may still not be operational. You may have excellent error performance on the cabling or transmission medium and still get a black picture. To verify the health of the MPEG multiplex being transmitted, the formats of the transport packets themselves must be checked. As stated in our earlier discussion of PSI, EPG, and conditional access information, many complicated data structures are transmitted in the multiplex and any, if corrupted, will lead to a service outage at the viewer premises. The DVB Project Measurement Group has defined a series of checks (in ETR290) that should be made to ensure the transmission of valid MPEG data.

These checks consist of a simple set of checksums and rule checks that indicate whether a valid multiplex exists on the received bitstream. These checks have been implemented in real time in the HP MPEGscope series of analyzers, and some are also being implemented by decoder IC manufacturers in demultiplexing chips.

The ETR290 checks are categorized in three ways:

- First priority: necessary for decodability. These checks establish a set of conditions that are needed for basic decoding. If the conditions are not met, then a serious service outage is occurring, and customers are not receiving service. The checks look for critical failures in synchronization, for proper PAT/PMT tables, for the existence of the PID streams, and for packet loss.
- Second priority: recommended for periodic and continuous monitoring. These checks look for conditions that indicate outages which may affect a portion of the services. Such errors are serious but do not necessarily indicate a complete service outage. In this category of checks are the verification of time base (PCR/PTS) accuracy, CRC checks on PSI/EPG information, and the presence of conditional access (scrambling keys) information.
- Third priority: application dependent monitoring. These checks look for conditions that are not as severe as the first two and that may only affect some services and viewers. In this category are checks of the PSI table formats, table repetition rates, and buffer overflow/underflow errors.

If passed, these checks offer a useful way to identify gross errors or equipment failures, but they do not guarantee that the MPEG data delivered is correct. That is to say, passing these tests will not guarantee correct operations, but will provide a good rough check to see that a valid MPEG signal exists on the link. Links should be monitored in real time on a continuous basis at key points in the network, so that when failures occur, they can be rectified quickly with a minimum of downtime. HP produces a piece of equipment, the HP MPEGscope DVB, that continuously checks for these conditions, logging any errors and sending message to a central network management console. (See Diagram 10.)

MPEGscope DVB ETR 290 Checks

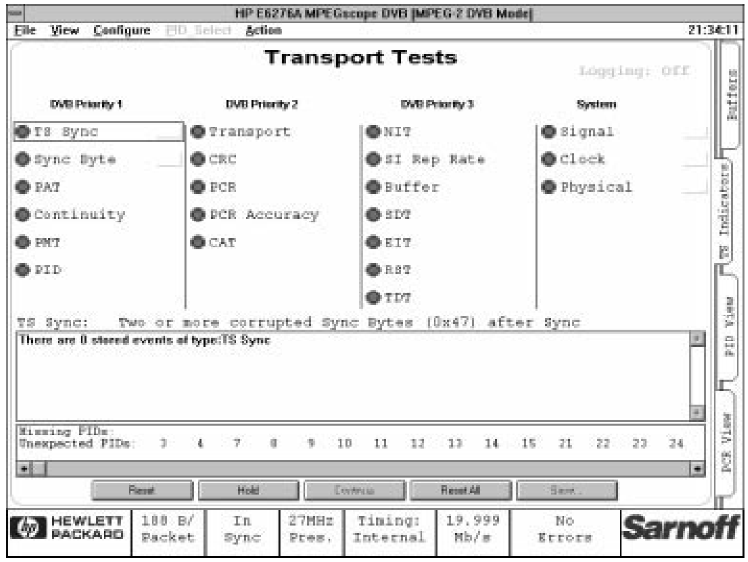

Typically, key points such as the master headend or satellite uplinks are prime candidates for continuous monitoring of the ETR290 health checks. At any of these points, the integrity of the entire equipment chain can be checked in one spot. In some cases, there may be many multiplexes, and to continuously monitor all the RF channels or transponders may not be practical. In such cases, some network operators are planning to install a single monitoring test device, and to use an RF switch or tuner to periodically send each multiplex to the test equipment for a fixed period of time. In this fashion each multiplex can be "sampled" and checks made for signal integrity periodically.

The checks defined in ETR290 are by no means comprehensive; they can be described as cursory at best, simply looking for gross errors. Nevertheless, they have the benefit of being simple enough to be performed on the fly, in real time, for every packet that is received by test instrumentation without incurring a lot of expense (our MPEGscope DVB uses a DSP card to implement this rule-checking). To fully verify the integrity of the MPEG transmissions, we must resort to offline tools.

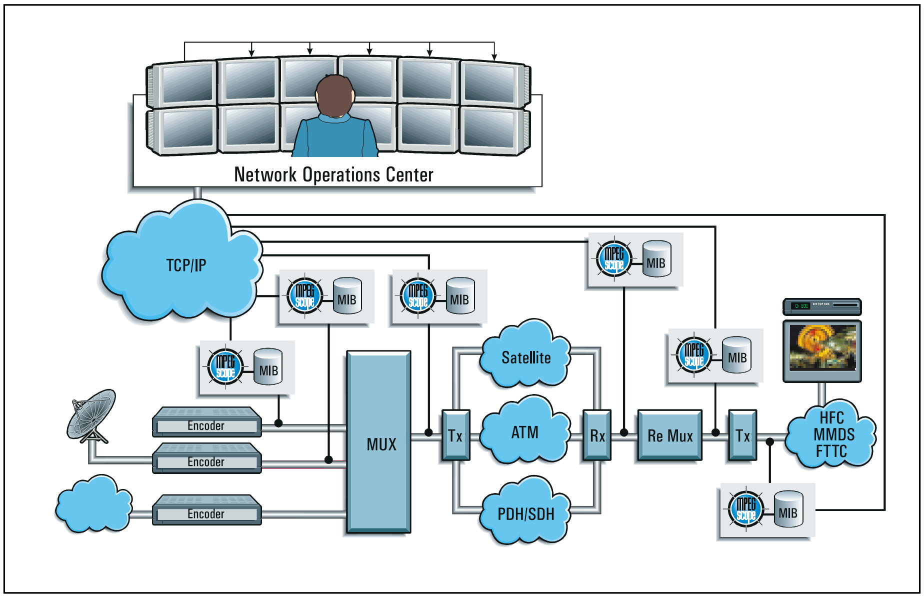

Diagram 10: MPEGscope Probe Monitoring System

If we capture a sample of the MPEG stream onto hard disk, then we can use more sophisticated rule-checking to completely verify the MPEG information. HP's DVB/MPEG compliance test software implements such a rule-checking system, which comprehensively verifies the format of the encoded streams and transmitted table information, and provides meaningful diagnostics about the discovered error conditions.

RF Measurements

The errors encountered in the digital MPEG transport domain are usually binary—a condition or error exists or it does not. But once the digital signal has been modulated into the RF domain, the errors become quantitative, and the accuracy of the modulation must be verified. While a complete treatment of RF modulation testing is beyond the scope of this paper, we will briefly examine how the modulation system works to identify any impairments it may bring to the digital signal.

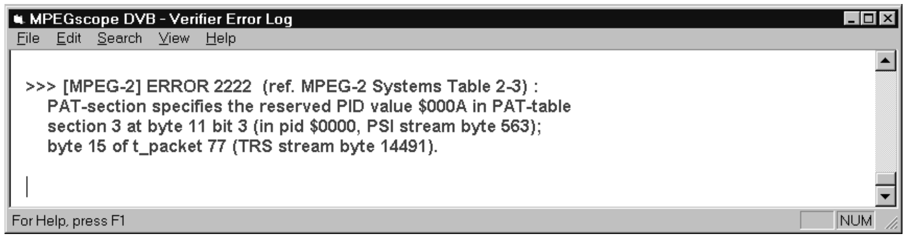

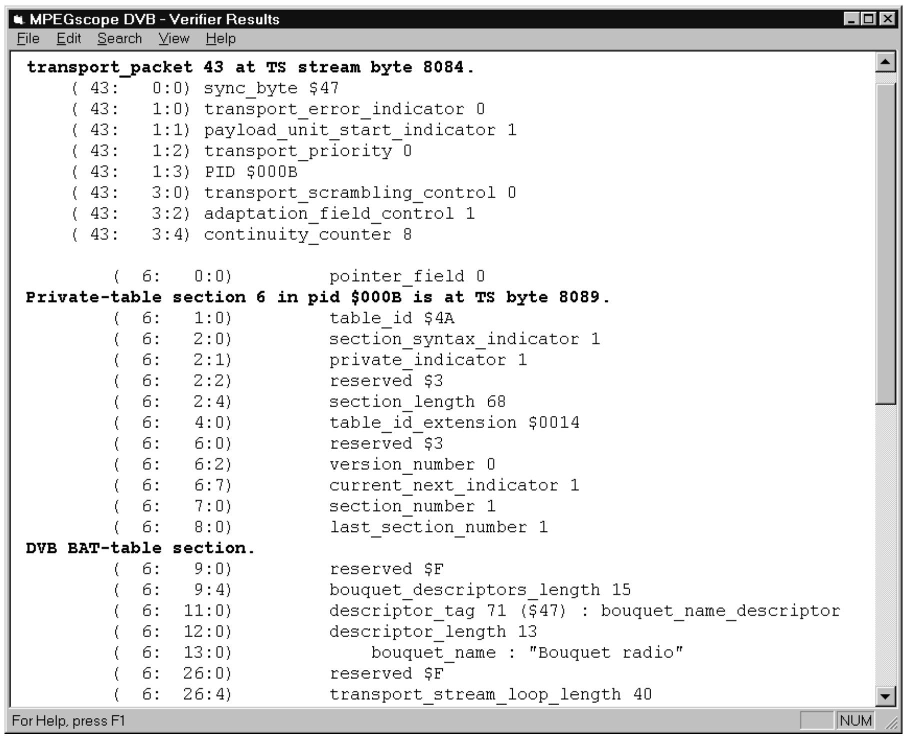

MPEGscope HP E6277A Compliance Verification Checks

The process of modulation translates a group of digital bits (called a symbol) into two dimensional RF information consisting of the amplitude and phase of the transmitted signal. The number of bits in each symbol varies with the particular modulation scheme; that is, 16 QPSK encodes 4 bits into each symbol, and 256 QAM encodes 8 bits in each symbol. Each dimension, (called in-phase I and quadrature Q) encodes a portion (e.g., half for QAM) of the bits in the symbol. For 64 QAM, the 6 bit symbols are encoded as 3 bits in the 8 phase states, and the other 3 bits in the 8 amplitude states.

The modulation can be verified in a number of ways:

- Visual examination of the modulation constellation. By plotting the received data points on a two dimensional I and Q graph, imperfections such as overdriven amplifiers can be detected by distortions of what should be a regular graph pattern.

- Numerical methods such as measurement of the deviation of the received symbol state locations (in terms of I and Q values) from the ideal state locations. These measures are called the modulation error ratio (MER) or the error magnitude vector (EVM).

- Traditional RF measurements such as the Signal To Noise Ratio (SNR), adjacent channel leak, and signal power measurement.

A more complete treatment of this subject can be found in our paper *Understanding the Measures of Signal Quality in DVB Systems.*

Errors in the modulation appear as symbols that have been incorrectly decoded at the demodulator, resulting in bit errors. The measurement of modulation quality and errors is a separate measurement domain from the test of the digital transport stream, because modulation impairments manifest themselves only as bit errors, and they can affect any of the transport stream bits; not just a certain kind of packet. Once we have verified that the transport stream has been correctly transmitted to the modulators, there is little point in verifying it again other than to check the bit error rate.

Modulation errors induced by either transmission noise or improperly calibrated equipment are compensated for by the forward error correction (FEC) system. But for any given link, the FEC system will compensate for errors up to a given level (designed into the system) only. If an error level exceeds the number of errors that can be corrected by the FEC design, then the system will fail dramatically. This leads to a behavior often dubbed the "cliff effect"—a step function in performance that occurs when errors exceed the critical level. When the error level is below that critical level for which the FEC can compensate, a link will seem relatively error free, even in the presence of a large number of errors. Then, all of a sudden, things may go drastically wrong–if the critical level is exceeded, the performance "falls off the cliff."

This effect underscores the importance of BER measurements in commissioning MPEG systems, because systems may work fine, even when their BER performance is marginal. The system may be close to the "cliff," and the slightest increase in noise, rain fade, or some other impairment will cause a catastrophic failure. This critical BER of the system is dependent on the network architecture.

System Specific Tests

Each of the current architectures being used for digital video has some unique system-specific issues that need to be considered for that architecture. The following are some key points that need to be observed in these kinds of systems.

Hybrid Fiber Coax

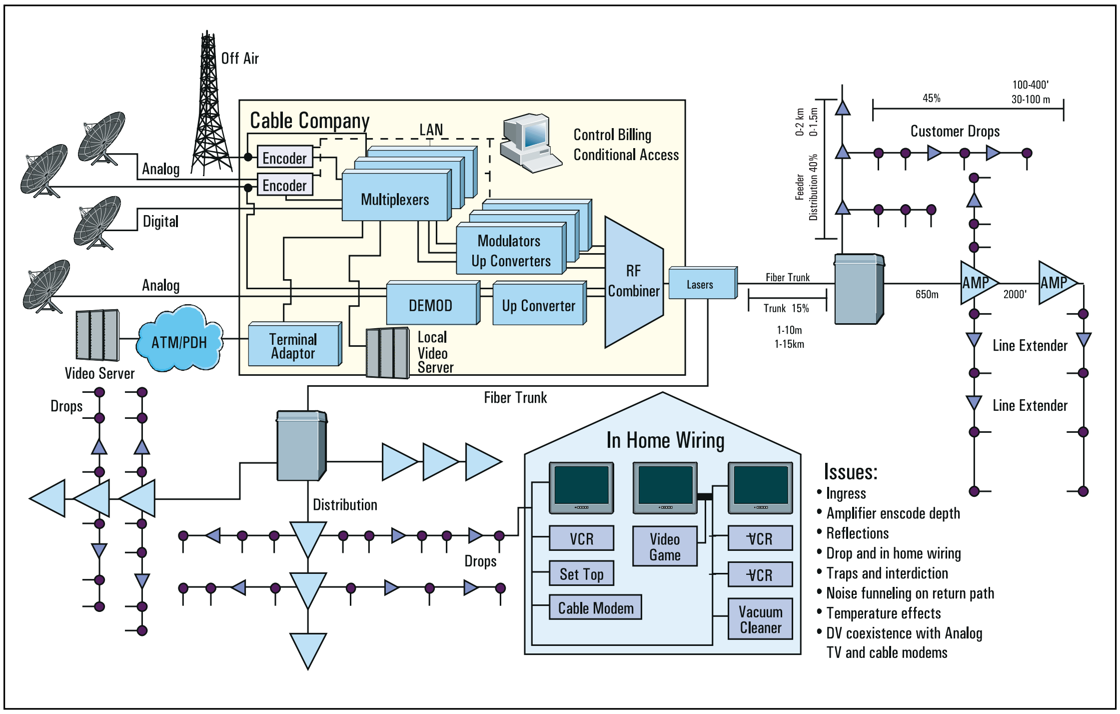

Diagram 11: Hybrid Fiber Coax

One of the most popular delivery mechanisms planned for DV is the existing cable infrastructure. There are two kinds of coaxial infrastructures that should be considered: new modern hybrid fiber coax (HFC) infrastructures created from scratch with DV in mind, and upgrades to existing cable plant.

Coaxially cabled TV has been around since 1950 (the first CATV system with wire pairs was built in 1948), when the Bob Tarlton built the first community antenna system using coaxial cable on utility poles in Pennsylvania, USA. The 1970s and satellite delivery of signals to the headend brought a boom in cable deployment and service. But the planners of these systems never envisioned two-way operations and digital applications. The major undertaking in this architecture has been to convert the pure coaxial architecture to the HFC architecture to extend the coverage distance and signal quality. This effort consists of replacing the coaxial trunks with fiber optics.

The major challenge in transmitting high bandwidth information over coax is overcoming noise and distortion buildup from the many distribution amplifiers that are required every 700 meters (2000 feet) on coax trunks and the 1 to 2 feeder line extender amps needed every 100 meters (300 feet) in the distribution portion of the network. This problem is more complicated in the digital system than in the analog system for two reasons: usually the digital system is more sensitive than analog video systems to noise and distortion, and the digital signals are often sent out at lower power levels than are the analog ones, in order to minimize adjacent channel interference from the digital modulation systems complete spectrum utilization (analog TV doesn't completely saturate the 6 or 8 MHz bandwidth as QAM modulation does).

These amplifiers are a problem particularly in the provisioning of the return channel, because so-called "noise funneling" in the tree and branch architecture is used. The amplifiers "funnel" all the ingress noise in the return path from the customer taps back to the head end. This noise can corrupt the return channel transmissions used for cable modems and control channels from the DV set-tops used for PPV and NVOD requests. Older systems may have as many as 50 or 60 amplifiers in cascade on a given path, and even upgraded systems with a lot of fiber plant can have amplifier cascades as high as 20 to 30, which makes distortion a critical factor.

A major challenge for existing infrastructures now is to locate and clean up the sources of noise in the system. Sometimes the only option in some systems is complete replacement with new equipment and cable. Unlike old legacy cable systems, new HFC infrastructures employ amplifier cascades as low as 2 to 4, as in the system Pacific Bell is deploying in California for video and telephony applications. In this system each fiber strand feeds a neighborhood node serving 120 homes (compared to 100 to 500 served off a single feeder in older systems) divided into four quadrant coax strands with 2 to 3 amplifiers on each strand.

Switched Digital Video

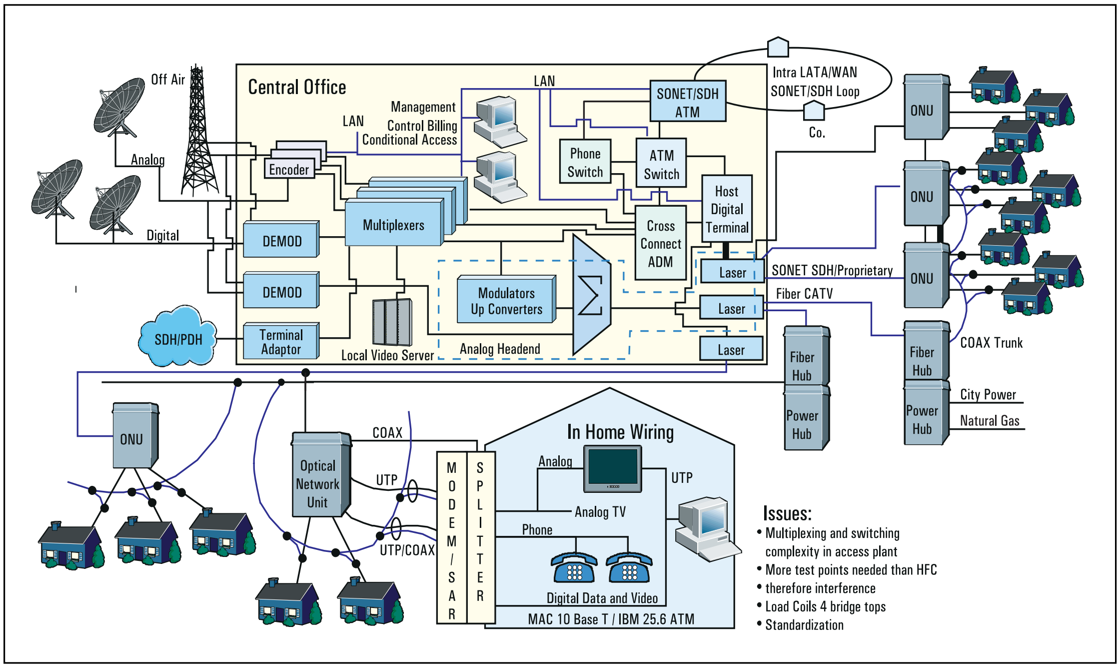

Diagram 12: Switched Digital Video

Switched digital video is perhaps the most complicated architecture being planned for DV—but it is also the most flexible in its ability to support future interactive applications. The complication in this architecture arises from the sophisticated equipment placed at the neighborhood nodes. Unlike the line extenders and amplifiers in HFC, which are passive to the signal contents, the SDV ONU pedestal equipment receives the fiber optic signal and performs multiplexing and switching functions to selectively transmit information to the star architecture, unshielded twisted pair strands or coax connecting the ONU to the homes.

This switching function adds complexity and flexibility. Complexity arises from the fact that the switching can malfunction. We need to test both the signal going in to the ONU and the signal to the drop, because they are different. This complexity, however, is the inherent power of the system—because it is switched, it can support flexible provisioning of applications from video to data and doesn't face the contention or congestion issues of a shared medium such as HFC or RF transmission. The bad news is that with switched architecture, simple RF tests out in the plant will no longer be sufficient, because there will no longer be one convenient monitoring point at the headend where all the digital signals can be tested and verified. The good news is deployment of more sophisticated access equipment will bring more capability for automated network management, diagnostics and testing, and potentially even greater network self-healing (a la SONET and SDH) and automated protection switching to back up facilities.

Other issues around SDV arise from limited standardization. There is a DAVIC SDV standard, but there are many vendor proprietary twists in the equipment proposed and on the market today. There are a variety of signal formats proposed on the fiber system and even greater variety in the transmission format used on the copper. On the copper side, the long-standing xDSL debate about which modulation scheme to use—the more readily available carrier-less amplitude phase (CAP) or the more powerful and potentially more expensive Discrete Multi-Tone (DMT)—continues vigorously in the standards bodies and the marketplace. The standards bodies are divided, with DAVIC choosing CAP and the ADSL Forum and ANSI choosing DMT.

Assuming that a solution comes from either the marketplace or the standards bodies, the other issue awaiting these systems is the old copper plant. The bandwidth supported and bit-rate available from these twisted pair modulation schemes decreases with the distance they must traverse over copper. It remains to be seen how close the ONU will have to be placed to the home and what will be the associated cost to use the extensive, but aging, copper plant to achieve the minimum 1.5 to 4 Mbps data rate needed to transfer a single broadcast MPEG stream. And in situations requiring two set-tops to be fed from an single drop, the rate will be an even more difficult 3 to 8 Mbps minimum. Only limited data is back from trials; undoubtedly we'll hear of situations in which the physical copper drop length will be an issue.

Other potential issues exist on lines with load coils or bridge taps, which need to be removed and replaced for these new modulation schemes. And finally, the high bandwidths used could pose a threat of significant cross-talk interference to other pairs cabled in the same sheath. The end result is that, as in the case of two-way HFC, significant reconditioning and rebuilding of the old plant could be needed to deploy SDV solutions.

Terrestrial Broadcasting

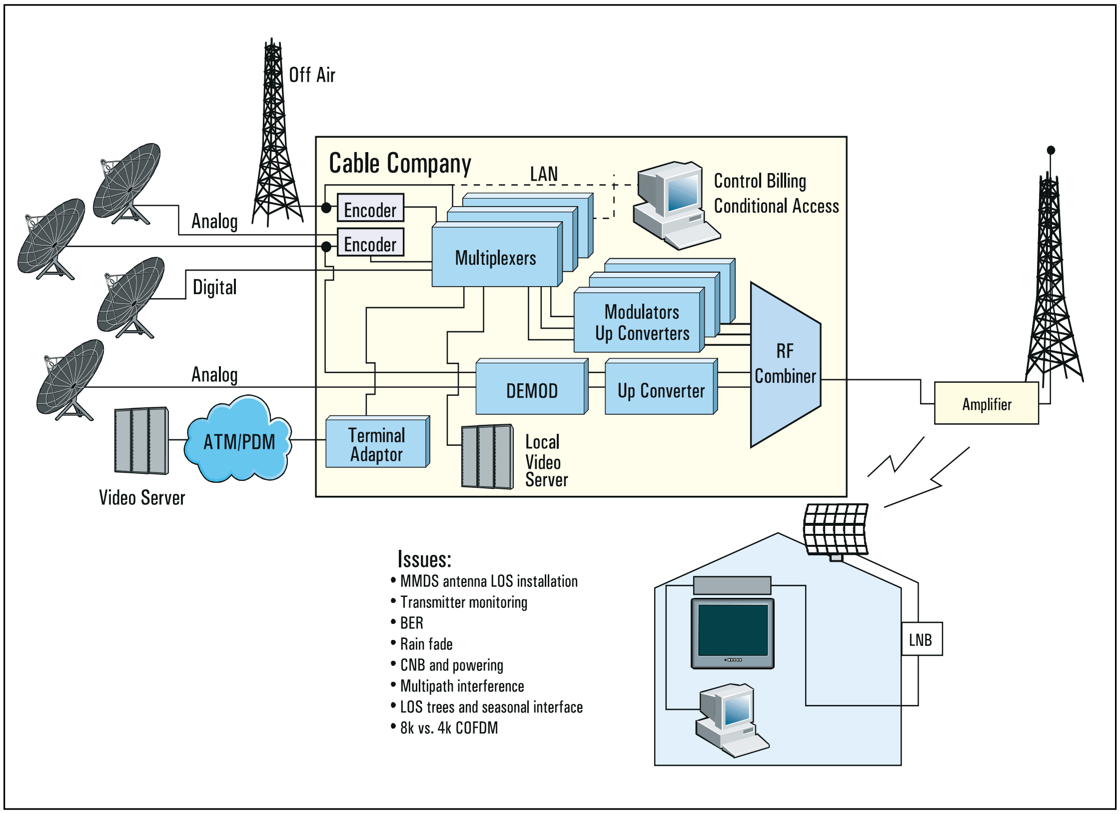

Diagram 13: Terrestrial Broadcasting

There are two kinds of terrestrial broadcasting being planned for DV: conventional TV style broadcasting using the DVB-T and ATSC standards, and microwave, line of sight MMDS or LMDS systems using multiple cellular-like transmitter installations. System -specific challenges for these systems centers around ensuring good signal reception at the antenna.

The large issues are signal strength and reflections during antenna location and aiming. Reflected signals, called "multi-path interference", pose a difficulty because they may lead to orientation of the antenna at a false transmitter—in reality just a signal bouncing off a building, or, even worse, some temporary object that will move, eventually making the reflected signal source invalid. If the reflections are close enough to each other to interfere, antenna relocation may be required. In both the DVB and ATSC terrestrial standards, the modulation schemes (COFDM and VSB) were selected because of their high multi-path interference resistance.

Some solutions to these problems are straight-forward, such as ensuring that installation technicians have a compass and a good map to help them align antennas to the real transmitter and not to a strong reflection. Other difficulties involving the antenna are physical, because unlike cable taps at the ground level, antenna installation usually involves climbing to a roof or other elevated location to get line of sight clearance. Installation usually involves orientation to maximize signal strength and reduce bit error rate, but it may be difficult for the person doing the installation to measure signal strength (and other things such as powering the low noise block on the antenna with test equipment) when he or she is hanging on for life from a precarious roof-top.

Weather and seasonal effects are also an issue here. We must ensure that the signal strength is adequate to compensate for rain fade during inclement weather. Another problem is leaves on trees—winter and fall installations that have marginal BER when no leaves are present may be pushed over the cliff when spring bloom arrives and leaves block the antenna.

But despite the difficulty of installation, terrestrial broadcasting has the distinct advantage of easy network maintenance. There is one convenient point, at the master transmitter feed point, where the signal can be monitored tested and verified. Other than checking the transmitter, there are no major system test issues. In multi-transmitter installations, the core network linking the transmitter sites needs to be tested and monitored as well, but this is a far cry from the complexity of maintaining an SDV or HFC plant.

Direct Broadcast Satellite

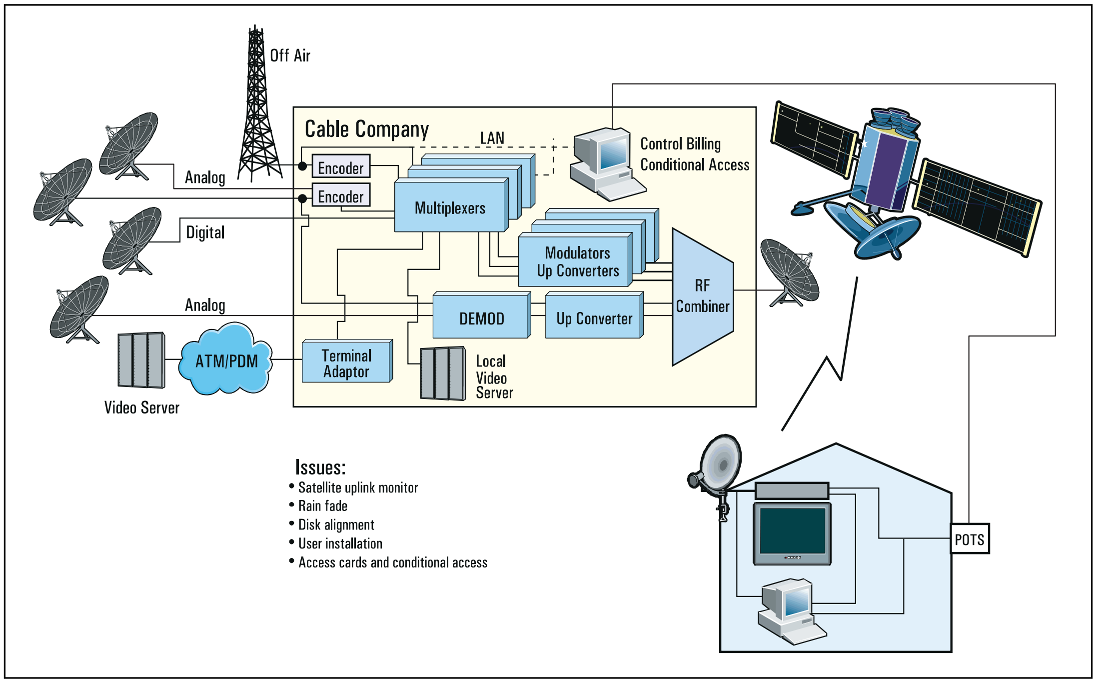

Diagram 14: Direct Broadcast Satellite

Like terrestrial broadcasting, digital broadcast satellite (DBS) is easy to test. Once we pass the enormous technological hurdle of putting up the satellite (and keeping it working there), the rest is almost simple. The key test point for this system is the uplink—after the signal at the uplink has been checked, there is little besides noise, interference, and rain-fade the will perturb the signal. The satellite doesn’t modify that uplink signal; it just re-broadcasts it. If there are no problems with the signal to the satellite, only dish orientation and alignment remain. Usually there is only one uplink and transmitter site, so the core network maintenance issues are straight-forward. The return channel for the system is usually the telephone network, where testing issues are well known.

The one area in which satellite systems are more complex is in the area of conditional access. The scrambling and billing system in a satellite service has to serve millions of consumers, as opposed to tens or hundreds of thousands of subscribers in the case of HFC, SDV, or MMDS. This means that the system is more crucial, carries a much higher volume of data, and is potentially more complex.

Asynchronous Transfer Mode

Unlike the solutions we have talked about up to now, ATM is not really an access and service delivery mechanism (with the exception of some SDV systems). ATM is a system usually used in the core network to link transmitter sites, to deliver information from remote servers to headends, to deliver and switch video streams in video on demand environments, and for long-haul transmission of MPEG video. ATM is complex and has some unique test issues that are worthy of exploration.

The big worry with ATM is the non-deterministic behavior of an ATM switch—cells don't always have the same delay going through the switch. Depending on the load on the switch, ATM cells spend variable amounts of time in the input and output buffers of the switching fabric. The fabric of the ATM switch is comprised of many small hardware switching elements that route the cells based on the virtual circuit number contained in the cell header. If contention for a path through the fabric exists, buffering will delay the cell until the path is free. If there is a lot of contention, the buffers may run out of room and the switch will be forced to discard the cell.

Diagram 15: ATM

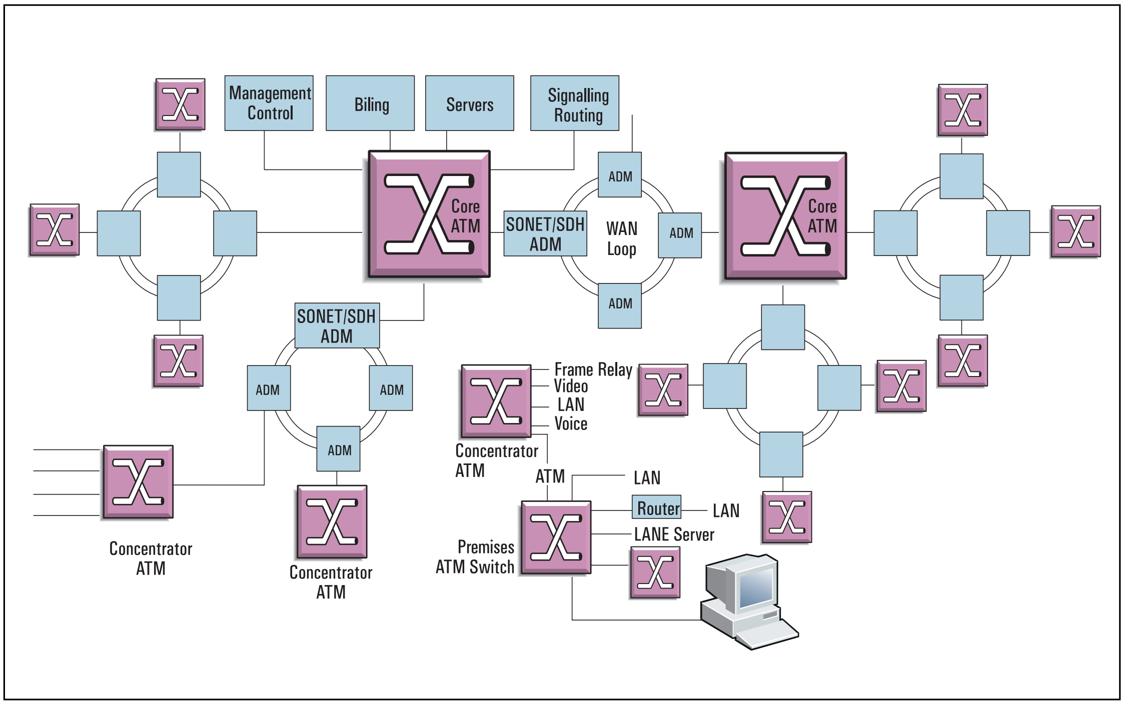

Diagram 16: Congestion, Cell Delay & Cell Loss

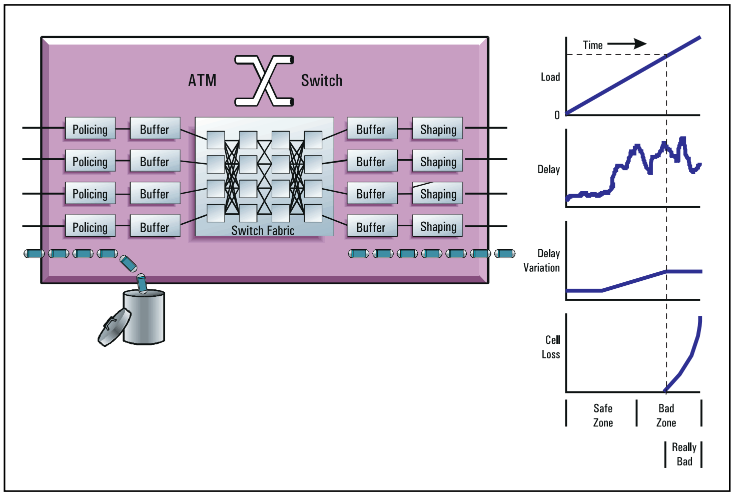

This process leads to a three state behavior in ATM switches. At lower levels of load the switch behavior will be very deterministic, showing nearly constant delay and no cell loss. As the load increases, the cell delay variation will increase, up to the point when the buffers become full. At this point the cell delay variation (CDV) remains unpredictable, but as the load on the switch goes higher, more cell loss will be experienced. The trick to building good ATM MPEG links is to make sure that the switch stays in the lower, stable region of operational load (which can still extend to load levels as high as 80% on some switches). The load level is enforced by policing hardware that rate-limits the inputs to the switch fabric. To ensure that the ATM switch stays clean and that it correctly transfers the timing-sensitive video information, it is important to configure the circuit bandwidths and policing hardware, either through signaling or PVC configuration, to make sure that the total load put on the switch doesn't push it into one of the "bad" areas of operation.

If the switch is pushed into a "bad" or loaded zone, the cell delay variation could introduce enough jitter in the PCR clock to make the received signal impossible to decode. Measuring PCR jitter is a key test in ATM environments, particularly if multi-hop ATM links are used, as the delay variation and jitter from multiple switches could be additive. Careful network design, load characterization of the CDV, and loss of the switch as load goes up are important to ensuring good system reliability.

Testing the Boundaries of the Network

We have discussed the transmission of the DV signals extensively. Let’s now turn our attention to the boundary of the network—the edges where analog video becomes digital and back again—the encoder and the decoder.

Decoder testing

Because MPEG decoders will be mass produced and are standardized, their operation is rigidly specified in the MPEG and related standards. This is not to say they are simple devices—because the opposite is true. For a consumer device, the modern digital set-top has unprecedented complexity. The computing power in the integrated receiving device (or IRD as the set-top decoder is sometimes called) rivals what a few years ago would be found only in high end scientific workstations.

The primary IRD test procedure is to activate the IRD, look for diagnostics from internal self-checks on the LEDs, and see if a picture comes up. If it doesn’t, swapping with another device should tell whether the signal is present or if the device itself has failed. Most vendors are putting special diagnostic modes into the IRD to help technicians troubleshoot. As we mentioned above, the set-top can provide a wealth of information about the received signal and the error rate being received.

In case of audio or video malfunctions, the decoding process in the IRD itself may be tested with a series of special test patterns. Researchers at the David Sarnoff Research Center have developed a unique test system that will exercise the myriad options built into the main MPEG decoding specifications. The MPEG audio and video specifications are very complicated and include many optional and not-often used features. The complexities can make identifying correct decoder operation sometimes difficult. There can be many subtle failures in the process of decoding the audio and video (or even flaws in the designs) that are difficult to detect with the eye and ear under normal circumstances as they flash by.

The clever part of the Sarnoff test sequences is that they are specially coded so that decoding flaws are propagated, from the short burst of data that exercises a particular option in the spec, into a series of frames following the test sequence. Each test consists of a preamble screen that lists the test being run; a short data burst that contains the test sequence; and a post-amble that contains screens with the word VERIFY written in large block letters. If the decoder isn't designed or working correctly, the decoding errors will be propagated into the VERIFY screen and show up as noise/garbage or some other visual effect. Using the test sequences, technicians and design engineers can verify visually the decoder operations—if the VERIFY screen comes up clean, then the test is passed. Recently the engineers at the Sarnoff Research Center, led by Dr. Michael Isnardi, have extended this methodology to audio decoding as well, using errors propagated into test tones to let technicians listen to the reference patterns in order to identify errors.

As complex and intricate as the audio and video processes are, the last and most complicated part of the decoder functions that must be tested is the EPG and return-path processing. Full coverage of these tests is beyond the scope of this paper, because the interactions and functions performed are so intricate. With the EPG alone there are many kinds of tables and information that need to be decoded. The essential tool for testing these functions is HP's DVB/MPEG compliance verification test suite, which enables the transmission of arbitrarily coded table downloads into the set-top. This analyzer is particularly useful because it is able to transmit broken, or incorrect tables as well as correct ones. One key point to test that is often overlooked in initial designs is the set-top's capability to reject and cope with incorrect tables. There is a high probability that operator error or a design bug may cause transmission of an incorrectly coded table. It isn't very user friendly to have the set-top lock up in this case, so it is essential that the EPG operations be verified with a number of input types and formats.

The return path transmissions go even one level higher in terms of complexity of interaction. They involve sending complicated data types (such as DSM-CC or Java applets) for ordering PPV or VOD movies or menus, and two-way protocol exchanges and state machines. Timing of the messages is as important as the formats of the messages themselves, because race conditions can occur. Timing is a complex area of digital transmission that promises to get even more complicated as more interactive and data functions are enabled in our video applications.

Encoder Testing

Testing the rigidly defined decoder audio and video functions is simpler than testing the encoder. The decoder is rigidly specified in the ISO MPEG spec, but the encoder isn't. As long as the encoder can feed its output to a standard decoder, it meets the spec. This fact has led to a wide divergence between encoder implementations and designs. Making it easier to select and test encoders until now was the fact that most of them were based on a handful of core chipset designs from companies such as C-Cube and IBM. 1997, however, ushers in a plethora of new encoder solutions, from single-chip encoders being produced by several manufacturers to single-processor DSP cores that are now fast enough to perform most of the encoding functions in software. Life is becoming very complicated indeed in encoder land.

One of the major issues in testing encoders is that their performance is heavily based on what input sequences are used. Images with lots of motion and high frequency content (for example, a tiny several-pixel region of major interest such as a tennis ball or hockey puck) are more difficult for an encoder to represent, particularly in cases in which the bit-rate is quite low. Different encoders have different bit budgeting and motion estimation algorithms that can lead to dramatically different performance on the same input sequence. Our HP Laboratories organization has been studying this problem extensively for many years, and we've concluded that this is indeed a tough problem!

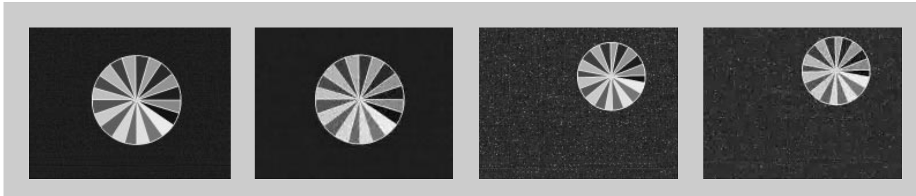

Diagram 17: Test Sequences

HP Laboratories has been working on test sequences to be used for evaluating encoders, and they have identified several pathological kinds of patterns that will torture-test most encoders. Some frames from one example are shown in Diagram 17. These sequences, in conjunction with the popular ISO committee sequences (flower garden, diva with noise, etc.), can be used to get some indication of the encoder quality.

Another problem is how to judge the output of the encoder. The major issue here is that the response of the human visual system to images is very complicated. The eye notices different levels of detail depending on whether the information is in the chrominance (color) or luminance (black and white) part of the image. The eye also is sensitive to different levels of detail depending on the motion of the subject and how complicated the image is, and whether the subject is in the foreground or the background of the image. In complicated images the eye tracks more rapidly over the subject, making fine details less relevant. The eye's response also varies from viewer to viewer.

In 1959, during development of analog television, a group called the Television Allocations Study Organization(TASO) in the U.S. studied the amount of noise, interference, and distortion that viewers would tolerate in a TV picture. The results were expressed in a five point scale rated excellent, fine, passable, marginal, and inferior. They eventually led to quantitative, objective measurement of tolerable distortion levels in reference test waveforms, which correlated a distortion level with a visual effect. These distortion levels are used as stringent quality of service guidelines for the licensing of commercial broadcasters. Unfortunately no DV equivalent currently exists.

Diagram 18: Subjective and Objective Testing

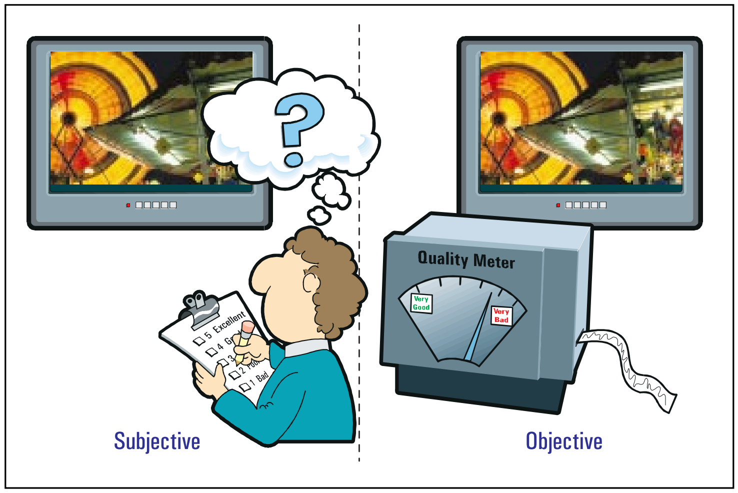

A standard, ITU-R BT.500 (see Diagram 18), exists for setting up subjective viewer tests using groups of human subjects to judge the encoder output (or TV set or whatever), allowing some sort of repeatability to DV encoder tests. Unfortunately, using this kind of psychological experiment as a test procedure is difficult to set up and repeat. Subjective tests like this are time-consuming, and what we desire more is some sort of objective or numerical method that will give us a quality scale similar to the T-pulse distortion masks used in analog TV.

One strategy for objective testing is to use a two-point measurement system in which the signal is tested at the input to the encoder and at the output of a reference decoder. The difference between these two signals is evaluated according to a model that approximates human visual response. The rub lies in the model—because researchers do not agree what model to use. Early tests used a simple mean squared error (MSE) comparison of pixel value deviation as a comparison model, but this proved to have very poor correlation to subjective test results. Recently the U.S. National Telecommunications and Information Administration Institute for Telecommunications Sciences (NTIA ITS), working in conjunction with ANSI, released a specification called ANSI T1.801.03 that describes a new model, which refines MSE, and that looks at peak SNR, changes in spatial frequencies, motion energy, and edge energy (horizontal and vertical contrast transitions, such as would be found in blockiness). The value of these tests is that they are computationally simple enough to be implemented in real-time hardware and used for operational monitoring and troubleshooting.

Unfortunately, the T1.801.03 model only achieves moderate correlation with subjective tests. A number of research laboratories such as Sarnoff have proprietary algorithms that have achieved a higher correlation of subjective tests to the human visual system. The difficulty with these more sophisticated models is that their computational complexity restricts their use to post-processing or to implementation in very expensive hardware, which precludes their use in operational monitoring.

The good news is that there is another option. Researchers at HP Laboratories have produced and applied for patents on a system that can identify MPEG encoder visual quality faults without having to look at the input sequence. The system uses what we call compressed domain quality analysis (CDQM). By careful analysis of the CDQM parameters and the contents of the MPEG encoder output, it is possible to identify when encoders are facing bit-starvation

Diagram 19: 2 Point Objective Quality Analysis

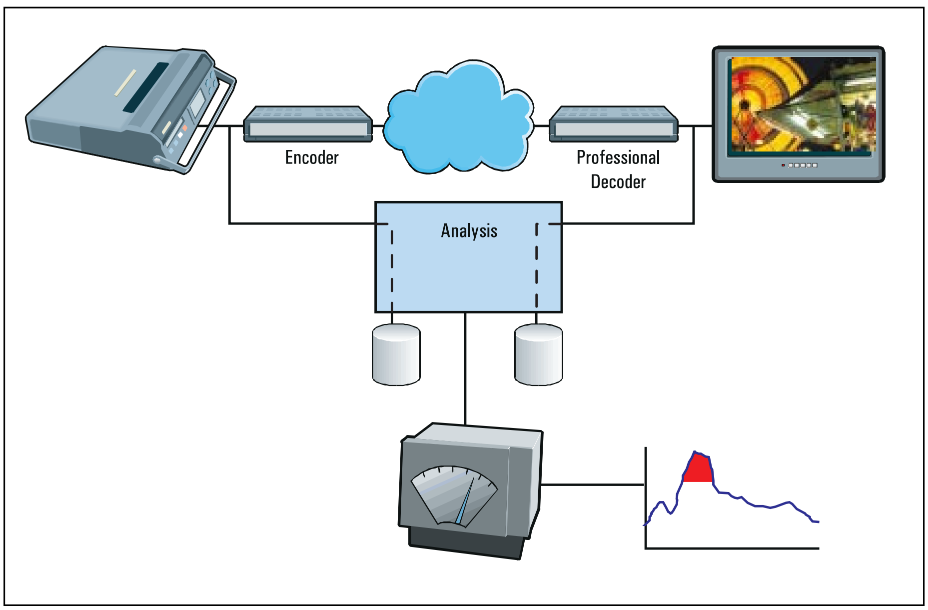

conditions that lead to excess quantization and introduction of visual artifacts into the original picture. This system has a number of advantages. First, it operates independent of the input image, so it requires connection only to one point in the network, anywhere on the compressed link. Second, it is simple and inexpensive to implement in low-cost hardware, allowing use in a wide range of continuous-monitoring, "encoder failure early warning" applications. The drawback is that CDQM works only with DCT-based MPEG compression. We will demonstrate this system at the National Association of Broadcasters show in 1997 and implement it on our MPEGscope series of analyzers.

This kind of test system will be crucial as the economic imperative to squeeze as many services as possible into the physical link will tempt operators and service managers to push the bit-rate slider on their encoder control panels down, to just get that “one more channel in.” Continuous monitoring will be needed to identify how often for the visual mix on a particular channel the encoder is introducing artifacts to see if the quality loss will be noticeable to viewers. The HP CDQM technology will fill an important need here.

Diagram 20: Compressed Domain Quality Analysis

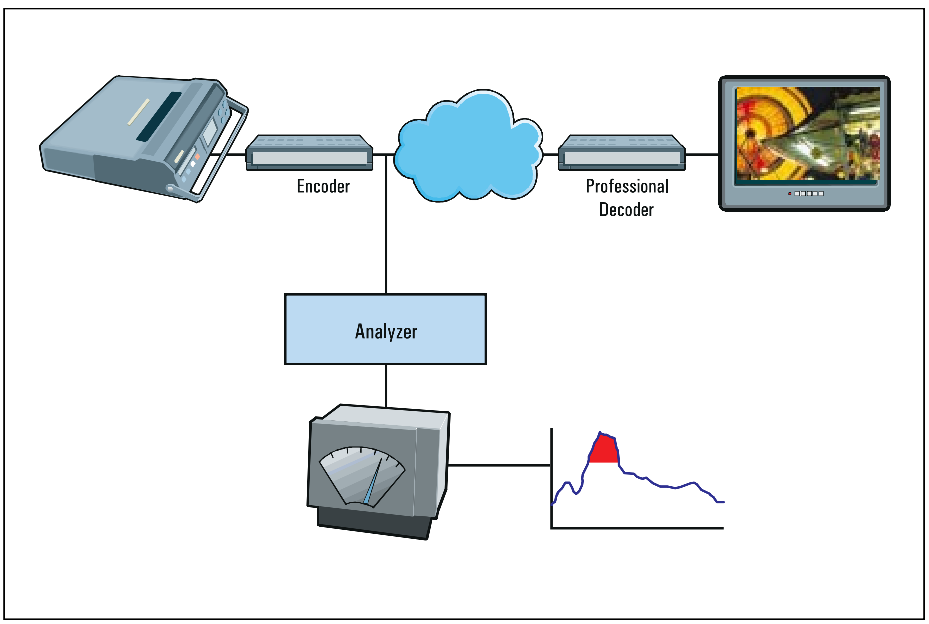

Commissioning a new DV system

As well encoders and decoders there are many components to be checked when assembling a new DV service, transmission links, encoders, multiplexers, servers and so on. It will be a complicated system, and there is much testing and characterization of the devices needed as homework before the system can go on the air. Unfortunately market pressures to sign up the subscribers run counter to the need for careful planning and technology testing. Mistakes tend to be more exponentially expensive to fix the further on in deployment they are discovered.

A balance will need to be struck between preparation and maintenance. At HP we are preparing test tools for both applications—in the lab before the system is installed, and maintenance tools for the field. Here are of some of the key products HP has introduced to assist the commissioning of these new DV systems.

- HP E4200 BSTS MPEGscope ATM
- HP J2306A Network Advisor
- HP E6276A MPEGscope DVB
- HP E6277A MPEGscope
- HP E4441 Reference Modulator
- HP ESG1000
- HP 89440
- HP 8594Q

For a further overview of Hewlett-Packard products, see page 161.

Glossary:

ADSL	Asymmetrical Digital Subscriber Loop
ANSI	American National Standards Institute
ATM	Asynchronous Transfer Mode
ATSC	Advanced Television Standardization Committee
BER	Bit Error Rate
BERT	BER Test
CAP	Carrier-less Amplitude Phase
CAT	Conditional Access Table
CATV	Cable Television
CDQM	Compressed Domain Quality Analysis
CDV	Cell Delay Variation
COFDM	Coded Orthogonal Frequency Division Multiplexing
CRC	Cyclical Redundancy Check-sum
DAVIC	Digital Audio Video Council
DBS	Digital Broadcast Satellite
DCT	Discrete Cosine Transform
DI	Discontinuity Indicator
DMT	Discrete Multi-Tone
DSM-CC	Digital Storage Media Command and Control
DSP	Digital Signal Processor
DSS	Digital Subscriber Satellite
DTS	Display Time-stamp
DV	Digital Video
DVB	Digital Video Broadcast
DVB-ASI	DVB Asynchronous Serial Interface
DVB-S	DVB Satellite
DVB-SPI	DVG Synchronous Parallel Interface
DVB-T	DVB Terrestrial
EMM	Entitlement Management Messages
EPG	Electronic Program Guide
EMV	Error Magnitude Vector
FCC	Federal Communications Commission
FEC	Forward Error Correction
HDTV	High Definition Television
HFC	Hybrid Fibre Coax
HP	Hewlett-Packard
IC	Integrated Circuit
IRD	Integrated Receiver Decoder
ISO	International Standards Organization

ITU	International Telecommunications Union
LAN	Local Area Network
LMDS	Local Multipoint Distribution Systems
MER	Modulation Error Ration
Mhz	Mega-Hertz
MMDS	Multipoint-Multichannel Distribution System
MPEG	Motion Picture Experts Group
MPTS	Multi-program Transport Stream
MSE	Mean Squared Error
NEM	Network Emulation Module
NTIA ITS	National Telecommunications and Information Administration Institute for Telecommunications Sciences
NTSC	National Television Standardization Committee
NVOD	Near Video On Demand
ONU	Optical Network Units
PAL	Phase Alternate Line
PAT	Program Association Table
PCR	Program Clock Reference
PES	Packetized Elementary Stream
PID	Packet Identifier
PMT	Program Map Tables
PPV	Pay Per View
PRBS	Pseudo Random Binary Sequence
PSI	Program System Information
PTS	Presentation Time-stamp
QAM	Quadrature Amplitude Modulation
RF	Radio Frequency
SNR	Signal to Noise Ratio
SPTS	Single Program Transport Streams
TEI	Transport Error Indicator
TS	Transport Stream
VLC	Variable Length Codes
VOD	Video On Demand
VSB	Vestigial Side-Band

HEWLETT PACKARD

For more information

For more information on Hewlett-Packard Test & Measurement products, publications or services, please call your local HP sales office. A current listing is available at: http://www.hp.com

United States:

Hewlett-Packard Company
Test and Measurement Organization
5301 Stevens Creek Blvd.
Building 51L-SC
Santa Clara, CA 95052-8059
1-800-452-4844

Canada:

Hewlett-Packard Canada Ltd.
5150 Spectrum Way
Mississauga, Ontario L4W 5G1
905-206-4725

Europe:

Hewlett-Packard
International Sales Europe
Geneva, Switzerland
+41-22-780-4111

Japan:

Hewlett-Packard Japan Ltd.
Measurement Assistance Center
9-1, Takakura-Cho, Hachioji-Shi
Tokyo 192, Japan
(81) 426-48-3860

Latin America:

Hewlett-Packard
Latin America Region Headquarters
5200 Blue Lagoon Drive, 9th Floor
Miami, Florida 33126 U.S.A.305-267-4245, 305-267-4220

Australia/New Zealand:

Hewlett-Packard Australia Ltd.
31-41 Joseph Street
Blackburn, Victoria 3130
Australia
131-347 Ext. 2902

Asia Pacific:

Hewlett-Packard Asia Pacific Ltd.
17-21/F Shell Tower, Time Square
1 Matheson Street, Causeway Bay
Hong Kong
(852) 2599-7070

5966-1032E 06/1997 Rev A
Specifications subject to change