SMPTE and VidTrans Joint Conference 2005

A Full Conference Report

"Moving Images Across the Screen and Across the World" was the theme for the first combined conference of SMPTE and VSF. SMPTE President Ed Hobson welcomed over 300 attendees and roughly 30 vendors to the symposium, held in Atlanta beginning Monday, January 31.

The first speaker to present was Richard Christen who gave an overview of video compression technology. MPEG-2 is a mature technology, having been in the industry for roughly ten years. Improvements in MPEG-2 coding have led to a reduction in bit rate of 50% to 75% compared to earlier implementations. Reasons for this reduction include better video preprocessing, better motion estimation, multipass encoding, and changes in GOP structure. Adaptive GOP structures allow encoders to vary the length of the group of pictures (GOP) and insert I frames based on scene changes. "Mixed frames" enable a single frame to contain I, B, and P elements. As compression has improved, MPEG-2 has reached a practical lower limit of 2 Mbits/sec. As a result, significant effort has been spent on improving MPEG-4. Improvements to advanced video coding (AVC) include lower overhead, improved motion vectors, and adaptive block size. VC-1 and AVC represent the future, but have different artifacts than MPEG-2 and offer challenges in the form of legacy support and in the development of robust support tools, which MPEG-2 now enjoys.

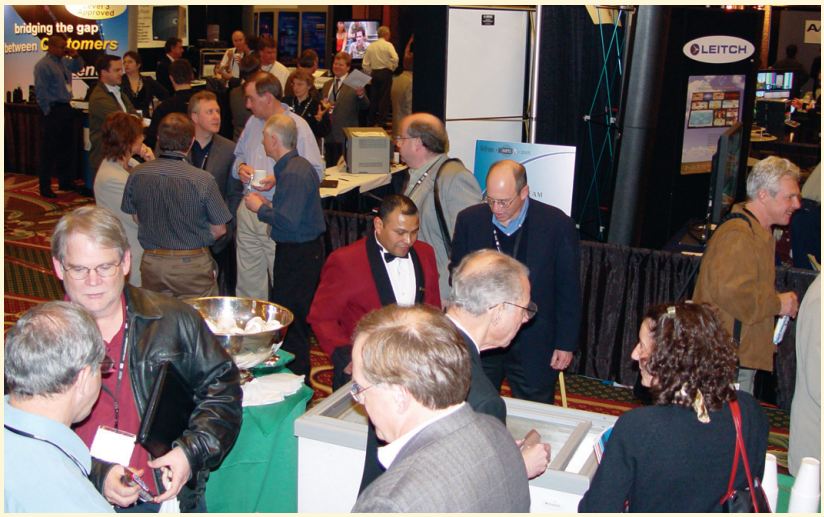

Next to present was Sridhar Srinivasan of Microsoft Corp., who offered an overview of VC-1 and began by making a distinction between VC-1 and Windows Media 9. VC-1 defines a bit stream while WM-9 is Microsoft's implementation of VC-1. VC-1 is an 8-bit compression system based on 4:2:0 sampling and is scalable in bit rate, making it suitable for a number of different applications and video formats. VC-1 transforms are based on (internal) 16 bit integers as opposed to floating point numbers and may be applied to block sizes of 4 x 4, 4 x 8, 8 x 4, and 8 x 8. This affords better reproduction of textures and film grains, and fewer ringing artifacts in the case of 4 x 4 blocks. In the case of motion compensation vectors, smaller blocks produce better predictions but more vectors to code. Subpixel resolution produces better resolution but creates more information to code as well. VC-1 supports three profiles and may be put into various containers such as MPEG-2 transport streams, as well as audio video interleave (AVI) and advanced systems format (ASF) wrappers. On the encoding side, deblocking filters reduce artifacts. On the decoding side, deblocking filters, deringing filters, frame resizers, and frame interpolators improved perceived visual quality. The VC-1 bit stream specification has been frozen pending SMPTE approval.

Paul Hansil spoke in lieu of Neil Brydon on delivering video at lower bit rates using advanced video coding (AVC) (MPEG-4 part 10). Similar to MPEG-2, AVC is a decoder-centric standard and defines the decode process and stream syntax. Four profiles are defined for a range of intended applications. The baseline profile supports mobile and conferencing, the extended profile supports streaming, the main profile for broadcast, and the high profile for broadcast and production. As with VC-1, integers replace floating point numbers in the discrete cosine transform (DCT) process. While decoder processes can be accomplished with presently existing integrated circuits, encoding processes are not presently done with mass-produced silicon. Therefore, a combined hardware and software approach makes sense at the present time. With high definition being up to six times the spatial resolution of standard definition, a different approach is needed to encode the higher bit rate material. One approach is spatial partitioning. With this approach, the image is divided into horizontal slices, and each slice is assigned a dedicated CPU. Much current research is in the area of Fidelity Rate EXTensions ("FRExt"). "Hi10" defines a 4:2:0 data structure with 10-bit samples; "Hi422" defines a 4:2:2 structure with 10-bit samples; and "Hi444" defines a 4:4:4 structure with 12-bit samples. AVC is on a steep improvement curve offering benefits in areas where MPEG-2 is challenged.

Michael W. Marcellin discussed JPEG2000 for digital cinema. Significant work has been done on JPEG, affording

many advantages over earlier versions. Significant advantages of progressive transmission (progressive in the sense of serially accumulating data as opposed to progressive raster scan) include lossless compression, user-selectable image size and resolution, as well as choice of region of interest (including aspect ratio). Two-color transform functions based on wavelets are supported. Irreversible color transform (ICT) involves conversion of RGB to YCbCr and is used for high-performance lossy compression. On the other hand, reversible color transform (RCT) is an integer approximation of ICT and is useful for lossy and lossless compression.

Wendy Aylsworth addressed the group in lieu of Mike Christmann regarding WorldScreen, a consortium dedicated to resolving system-related issues affecting worldwide adoption of digital cinema. WorldScreen currently faces two major challenges. First, the DCI specification only deals with the backend of the chain (from the distribution master to presentation in the theater; not the production process). Second, an optimized workflow model for the complete production chain for digital cinema in 4k does not exist. Consequently, the consortium is investigating a layered scheme compression technology to use throughout the production process. Secondary challenges involve implementing a layered compression scheme for "rich media" archives and integrating metadata into the production process that will survive through archiving.

Chris Purdy broached the subject of advanced compression system testing. The face of compression is changing. High-definition broadcasting is slated to begin in the U.K. in 2005 via satellite. IP delivery is becoming more common. Even the next generation of DVDs will not be based on MPEG-2. There are major differences when AVC is compared to traditional MPEG compression. Non-existent in AVC are layers, scenes, objects, and sprites. A frame can reference up to five other frames. Gone is the specific distinction of I, B, or P frame. A frame can contain slices of all three types. All of this requires new analytical tools. New tools spot non-compliance in bit stream syntax, inefficient coding, artifacts introduced by codec errors, and video buffer verifier (VBV) errors. New tools are capable of displaying macroblock types via color coding as well as displaying motion vectors graphically. A new MPEG-2 transport stream analyzer identifies advanced coding for elementary streams, showing levels and profiles, as well as video coding parameter descriptors. New test and measurement tools will enable engineers to design and maintain compression and distribution systems more efficiently.

Charles Ganzhorn began the Monday afternoon sessions by providing Part 1 of Video over IP. Growth in a worldwide infrastructure has led to a proliferation of cheap IP bandwidth. Digital TV is a strong driving force as is the transition from standard- to high-definition in creating a move to IP transport. Video and voice are the most profitable services offered by providers, who view this as "their" market. Consequently, the service industry is moving toward "triple play" (a network that supports voice, video, and data). Recommendations to accomplish triple play involve network characteristics of very low jitter (less than 50 msec), a bit error rate of 10-7, and end-to-end latency equal to or less than 1 sec. Quality of Service (QoS) elements involve traffic management and security. Network traffic control involves classifying data, placing it in a queue, and then shaping and fragmenting it. A number of threats exist from hackers, however, a number of tools are available to protect against threats.

Part 2 of Video over IP followed with a look at multicast technology by Beau Williamson. Advantages of multicasting include enhanced efficiency due to reduced server output, improved network performance due to the elimination of redundant traffic, and expanded functionality as a result of more services offered. The basic premise of multicasting is that no end user wants to receive a stream unless they specifically ask for it. Internet Group Management Protocol (IGMP) version 2 defines how a user asks the network for a desired stream and how it terminates. Distribution may be accomplished by two methods. In the "Source" or "Shortest

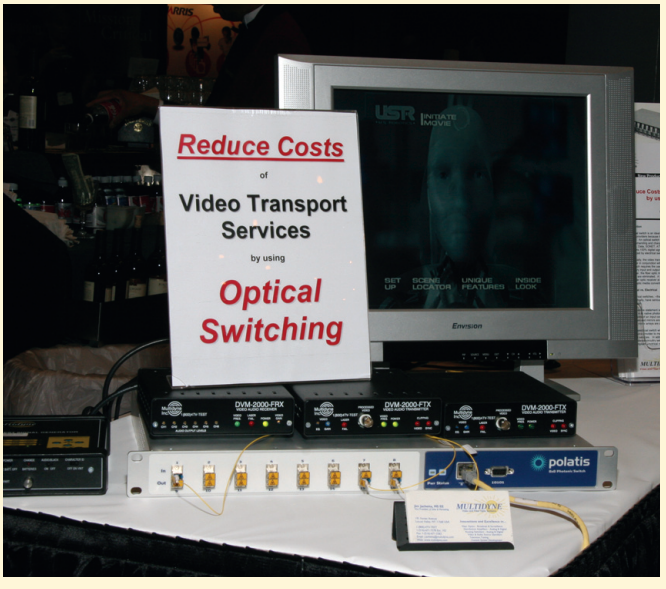

Path Tree" model, each source finds its own route to the receivers. In contrast, a "Shared Tree" model uses a "Rendezvous Point" to merge multiple streams that then take a common path. Rendezvous points may be created manually by network configuration, or may be dynamically assigned. A method known as "Anycast RP" assigns multiple RPs for the same group and allows sources and users to select the closest RP. Greater redundancy is inherent in this model.

Discussing "Managing Diverse Media Transport Assets," Anthony Magliocco outlined potential physical resources, the services possible, and the technologies used to provide said services. Managing media transport assets involve elements of administration, security, and monitoring. Administration involves bandwidth scheduling, billing for services, and device control. Potential devices to be controlled include data switchers, video routers, encoders, decoders, and other devices. Monitoring involves QoS as well as collecting data for long-term analysis to insure efficient bandwidth utilization over time.

Olaf Nielsen presented "Essential Elements of Triple Play." Today's current data service offer flat bandwidth at a flat-rate price with no priority afforded to disparate data types. Current phone service is protected from data, thereby insuring a guaranteed bandwidth, and video service allocates fixed bandwidth for broadcast TV and a fixed aggregate bandwidth for video on demand. However, dynamic service would optimize network capacity and revenues. That is the goal of the TR-059 Triple Play model. What is needed is classification of streams and deep packet inspection to identify dynamic applications in order to allocate bandwidth on the fly.

Tuesday's sessions began with John Dale offering background information on VSF and the challenge of updating standards of QoS and network reliability to cover digital and compressed signals. Brad Gilmer continued the theme and related the old paradigm to the new. Older network availability specifications were based on percentage (99.9%, 99.99%, 99.999%) of network "uptime." Current specifications are based on the number of network events occurring within a specified time frame, where a network event is defined as a single occurrence of packet(s) loss in the network. With bit rates of 1 to 400 Mbits/sec, service quality categories are less than .1 event per day, less than 1 per day, and less than 10 per day. Forward error correction (FEC) can help correct packet errors, but require additional bandwidth. FEC also increases latency, which may not be desirable in broadcast applications. Network errors can also be bursty in nature. By adding interleaving to the data, burst errors can be mitigated at the expense of latency. Dale continued with some basic characteristics of all packet networks. Specifically, errors occur in every network, IP delivery drops errored packets, adds jitter to packets, and delivers packets out of order. In order to quantify network quality, VSF, working with ITU, has developed recommendation Y.1540. These metrics compare a number of measurable parameters and relate them to system performance.

In the next presentation, Pierre Costa stated that in order to facilitate a diverse group of requirements, networks are first divided into unconditioned and conditioned service. In the case of conditioned service, the following is needed: FEC must support replacing lost IP packets, IP packets must include timestamps to aid in IP packet jitter removal, and a packet order number must be included to reorder out-of-order packets. Three profiles are also proposed: contribution, primary distribution, and access distribution.

Jeffrey Jensen pointed out that compression magnifies packet errors. Bit error rates (BER) for today's networks typically vary from 10-5 to 10-7. Unfortunately, broadcast requirements vary from 10-8 to 10-11, making FEC necessary. Of the various methods available, the most promising appears to be COP3 (Code of Practice), developed by the Pro-MPEG Forum. If packet data was placed into a matrix, FEC calculations could be applied to either vertical columns (original method) or to vertical columns and horizontal rows (COP3 Issue 2). Through choice of method and the size of the rows and columns, a tradeoff can be achieved between overhead and BER improvement. BER improvement varies between a factor of 10 and 100,000 just by using release 1 of COP3.

Pierre Costa offered real-world network experience based on a challenge to design a fiber-to-the-neighborhood, DSL-to-the-home IP network. This network would feature super high-speed data, high-quality video, and voice over IP. Subjective and objective testing determined the required video quality level (and therefore video encoding rate). The preliminary results of error analysis corroborated the theoretical modeling and simulation. Results indicate that FEC and interleaving is critical, latency may be an issue, characterizing network errors is essential, and end-to-end testing should be performed.

Danny Wilson discussed the topic, "Triple Play—Be Careful What You Ask For." He mentioned that the modern network has lots of "adjustments" to tune the network. There are now more channels than ever before. Customer expectations have grown, as well as the challenges to providers; namely that the internet is full of worms, viruses, Trojan horses, and additional traffic from spam and teenagers trading music. What is needed is a new suite of tools to effectively manage a network. These tools need to look at a number of disparate parameters, which could include RF spectrums, encoding and multiplexing parameters, transport streams, and IP statistics, to name a few.

Tuesday's afternoon session began with an overview of changes in newsgathering, by Howard Meiseles. Earlier, the change from film to videotape, and more recently the change from analog to digital and file-based. Newer technologies such as videophones, laptop editors, and digital satellite newsgathering push the envelope even further. In January 2003, IEEE passed an amendment to the 802 standard for the creation of wireless metropolitan area networks (MANs). "WiMAX" is the common name given to 802.16a. 3G is a new wireless standard for cellular service providing a payload of up 2 to Mbits/sec.

Rick Harding spoke on why the use of digital intermediates (DI) is changing the way motion pictures are made. In 2002, less than 3% of the major releases were produced using digital intermediates. Industry observers estimate that by 2006, 80% of all major motion pictures will be posted with digital intermediates. Currently, 2048 x 1556 pixels is the *de facto* standard for DI. From this, all requested screen formats may be derived. Because most screen formats are less than 1080 lines, this opens the door for a 1920 x 1080 10-bit format to be used as a digital intermediate.

Next, Harald Riner gave details of a new, high-definition, high frame rate camera. In the past, film cameras served this niche. Film cameras, however, have limited recording times, expenses, and delays associated with film processing, and other negatives that make video attractive. The new camera features a single CMOS sensor with a pixel count of 1536 x 1024. Refresh cycles place charges on each pixel that light proportionally discharges, thereby reducing vertical smear. Images may be captured at up to 1000 frames/sec or higher if smaller spatial resolution is used. The camera stores motion images on DRAM internally.

Hans Hoffman reported on HDTV from the European perspective. By not adopting first- and second-generation HD equipment, Europe has the luxury of a well-thought approach to high-definition and third-generation equipment. The European Broadcast Union (EBU) has decided on 720p 50 as their emission standard. They believe that interlaced was a good choice in the past for bandwidth reduction, but progressive scan compresses better at low bit rates, is more IT friendly, and offers better motion portrayal for sporting events. Production should be done in 1080p 50, and digital film should be done in 1080p 25. Present research is being done on developing guidelines for a mezzanine compression level. Initial thoughts are that horizontal subsampling is not good, and a bit rate greater than 140 Mbits/sec is indicated. To maximize the viewing experience, criteria for consumer display devices will also be established.

Innovation in HD studio lens design was the next topic, presented by Larry Thorpe. The design objectives were to create a new studio lens that was significantly smaller in total volume and much lower in weight while still maintaining full studio performance. The lens should have no compromises in operational capabilities or in interfaces to the outside world. After much research and computer simulation, such a lens was developed. At roughly 28% of the weight and 27% of the volume, its ergonomics are very good and performance is excellent. Next, the topic of lens performance and specifications was broached. Modulation transfer function (MTF) correlates spatial frequencies with associated amplitude and varies across the image plane. It also varies with focal length and aperture setting. Other things to consider when evaluating a lens are color reproduction and the concatenation effects of the lens, prism, and camera electronics.

The final speaker for Tuesday was Paul Chapman, who discussed emerging workflows in HD production. Hollywood is changing from a tape-based paradigm with linear editing to a file-based paradigm with nonlinear editing. "Data dailies" and "virtual telecine" are common buzzwords, as is "digital intermediate." While digital intermediate started as a scan of cut negatives to affect color correction, it now involves editing as well. One challenge is keeping track of all the source media. Time code is difficult when working with files.

Standardization in file types is needed. Color correcting with data files is trickier than with conventional sources and a universal method to tie color correction to the time line needs to be implemented. The goals of virtual telecine are to detach the image conversion from other processes normally performed at the telecine stage, and to produce a universal file from which different output formats may be derived. Allowing cropping and pan and scan after the transfer process will improve workflow. Similarly, color correction would be done at output formatting time as well. Digital dailies involves integrating metadata with separate audio and video files on to a timeline for daily review. What remains to be done is continued development and enhancement of emerging standards such as AAF, MXF, DPX, and the American Society of Cinematographer's proposed "Color Decision List."

An ice cream social followed Tuesday's sessions. Two chartered busses with over 100 attendees departed for a tour of selected parts of Turner Broadcasting's Techwood campus. Areas visited included Turner studios, post-production facilities, and the broadcast operations center.

The third and final day of the conference began with Bob Edge's presentation on building products and production systems with MXF. Much work has been done to develop standards for MXF, and a lot of standards have been created. A few items have surfaced in the area of interoperability and are being addressed by the standards committee. These involve how time code is handled, the use of ancillary packets, VBI data, and the challenges involved in using 10-bit data in 8-bit infrastructures. Engineers designing systems should ask vendors about cross-vendor support or perform empirical testing before committing to purchase equipment. Much work will be done in 2005 to resolves these discrepancies, and MXF will improve and evolve over time.

Continuing the MXF theme, Joe Fabiano spoke on how his company was remapping data to marry the MXF and AAF formats to XML data in order to achieve a hands-free system of content distribution for syndicated shows, news items, and commercial ads. Supported metadata includes time code, closed captioning, decoding parameters, and format sheet information (text). All objects are distributed in a standard IT packaging format, and by using an MXF implementation they are attempting to reduce the integration burden and maintain open standards. The next API will offer traffic and billing systems access to data presently going just to automation control.

"Achieving the Four Rights: Getting the Right Content to the Right Subscriber at the Right Time on the Right Device" was the topic of Norm Slater's presentation. Over 85% of available rich media is not re-usable. There is increased demand for content from more viewers and across a global network. What is needed is better tools to classify, search, and store data in a way that will allow it to be repurposed. The processes involved include content editing, metadata extraction or creation, library indexing, content management and analysis, rights management, and distribution.

Chris Parkerson spoke more about digital rights management. Balancing consumer expectations and content

providers' rights is a challenging task. Consumers expect ease of use, portability between devices, and often feel that if they've bought it they own it outright. They do not want to see visible signs of DRM. By contrast, content providers are concerned with revenue capture, property protection, and limiting redistribution after content is delerived to consumers. A good DRM system should ensure only authorized access to the content and implement restrictions necessary to enforce copyright protection in ways that are acceptable to the consumer.

Merrill Weiss spoke in lieu of Patrick Attallah on "ISAN"—an international standard for audio/video numbering. ISAN is an ISO standard that provides a unique, registered identifier for each movie and television show. The ISAN number would be the unique identifier and key to related databases. In the case of a movie, for example, it could link to other databases on the cast and crew, residual obligations, production details, and other metadata. An ISAN number would consist of a 48-bit root number, a 16-bit episode number, and a 32-bit version number that would be expressed in hexadecimal. Version numbers could distinguish different releases based on languages, screen format, "edited for TV" flags, and so forth. For a nominal cost, content may be registered through www.isan.org.

Wendy Alysworth presented on the broadcast flag and system renewability messages. The broadcast flag is a control descriptor consisting of a key, length, and value (KLV) triplet that is inserted in an ATSC stream in either the program map table (PMT) or the event information table (EIT). When present, it means "technological control of consumer redistribution is signaled." If the PMT/EIT is not inspected, content is considered unscreened. If the PMT/EIT is inspected, content is determined to be either unmarked or marked. Unmarked content may be sent anywhere. Marked or unscreened content may be passed to analog outputs, DVI at 720 x 480 at 30 frames/sec, a protected authorized digital

output, or an authorized-method recording device. A protected digital output could be HDCP (a protected protocol) traveling over a DVI cable, or DTCP (a protected protocol) traveling over IEEE-1394. System renewability messages is a methodology whereby media will not play in a compromised device.

After lunch, Andrew Karlin addressed the group on managed interfine transfer (IT) for television news operations. Television news is a "race for eyeballs." The challenge is how to get video faster and cheaper. Additionally, getting news from a local live truck back to a station and then to a national network is a slow, labor-intensive, manual process with no path for metadata. In the case of "soft" (feature) news, 90% is nonrealtime, but news networks still pay realtime rates. What network news producers want is immediate access to local news on demand, metadata along with news packages, fewer manual processes, the ability to repurpose content, and less dependency on videotape. All this suggests IT-based solutions. IT now offers more bandwidth to more places at less cost.

Wes Simpson presented test results from uncompressed transport of HD signals over IP networks. First, why use uncompressed signals? Benefits to broadcasters include very low delays (milliseconds compared to frames for MPEG compression), lack of compression artifacts, and support for embedded signals such as audio, time code, and metadata. Benefits to the carrier include simple setup and configuration, and support for a wide variety of video formats. Actual implementation was via IP over SONET operating at OC-48 speeds. Reed-Solomon error correction was used, as well as packet interleaving, to further reduce errors. Various test setups were configured to assess the effects of additional network traffic, and various routing schemes were compared. The conclusion: uncompressed HD video can be successfully transported over IP networks. A buffer is required inside the IP network adapter to remove jitter and meet SD and HD specifications. Video packets must be given priority over other data traffic in the event of competing traffic on the network.

Robert Castellano spoke on resilient packet ring flow control to achieve full convergence of triple play on residential and commercial networks. The requirements for convergence are multimegabit service with burstable bandwidth, high network availability, and ease of provisioning. The new goals are dynamic bandwidth reclamation and guaranteed QoS per subscriber. Traditional packet networks based on store and forward are inefficient. If a network is properly designed, data buffering can be eliminated if peak rate flows never exceed node capacity. Data entering the network should be classified and regulated by flow-shaping. Flow control should have closed-loop feedback to enable dynamic resizing based on available network bandwidth. This way the entire network capacity is used.

Next, May-Luc Champel presented on securing video contribution and primary distribution over IP networks. Thanks to higher bandwidth, IP networks now offer opportunities for transporting video streams. The Internet Engineering Task Force has produced several standards for defining video transport over IP. Bit alterations at the physical layer are converted into dropped frames by higher layer CRC mechanisms. Packets may also be dropped if network traffic becomes too congested. The Pro-MPEG Forum's COP3 FEC is designed for video over IP and is implemented at the network transport layer. Unlike other FEC schemes, COP3 relies on simple algorithms that can be implemented without hardware-specific parts. COP3 is adept at correcting single packet or burst packet losses and can be adapted to many different types of network conditions due to the variety of adjustable parameters. Lastly, it is transparent for devices that do not support it.

In the final presentation, Richard Bullock presented a case study on contribution quality video over IP. A "quality triangle" exists with the three corners being robustness, overhead, and latency. The key concerns in IP delivery of video are packet jitter, packet drops, network outages, and out-of-order packets. Ideally, packet jitter should be less than 40 msec. Multicast is becoming the desired technology. The case study involves the RTL Group, based in Luxembourg, and is one of Germany's major broadcasters. It has production facilities throughout Europe and North America. In early 2004, the contribution and distributions system converted to MPEG over IP. Multicast feeds leave master control at 8 Mbits/sec headed in eight directions. The network is comprised of leased lines. The IP routers were configured and are operated by RTL, which developed the IP management system as well. In the worst case, there are two network events (packet drops) per hour; they use COP3 issue 2. RTL is convinced that IP distribution is viable, and have more developments in the pipeline—such as advanced coding, centralized management, and integrated monitoring.

After questions and answers, the conference came to a close. The next SMPTE Technical Conference and Exhibition will be held November 9 to 12 in New York City; plan to attend!

SMPTE and VidTrans Joint Conference, ProMPEG COP3R2 Interoperability Demonstration

The Video Services Forum (VSF) is a trade association dedicated to addressing issues surrounding delivery of video over terrestrial networks. The VSF has several working groups. One of these groups, the Professional Realtime Video/Audio over IP Network Infrastructure Working Group, planned and conducted a successful demonstration during the Atlanta conference. The demonstration illustrated two things: (1) the transport of high-quality realtime video over a packet network with deterministic packet loss, and (2) interoperability between multiple vendors using independent implementations of the Pro-MPEG Forum Code of Practice 3 Revision 2 (COP3R2).

During the demonstration, each participating company generated an IP multicast consisting of a single 20 Mbit video stream with predetermined FEC/Interleaving. Prior to being distributed to all participants, the multicast streams were fed into a device (IneoQuest Technologies Singulus G1-T), which introduced deterministic errors. The errored streams were then distributed over the network to all other participants. This was done to demonstrate the recovery capabilities of the forward error correction (FEC) system described in COP3R2.

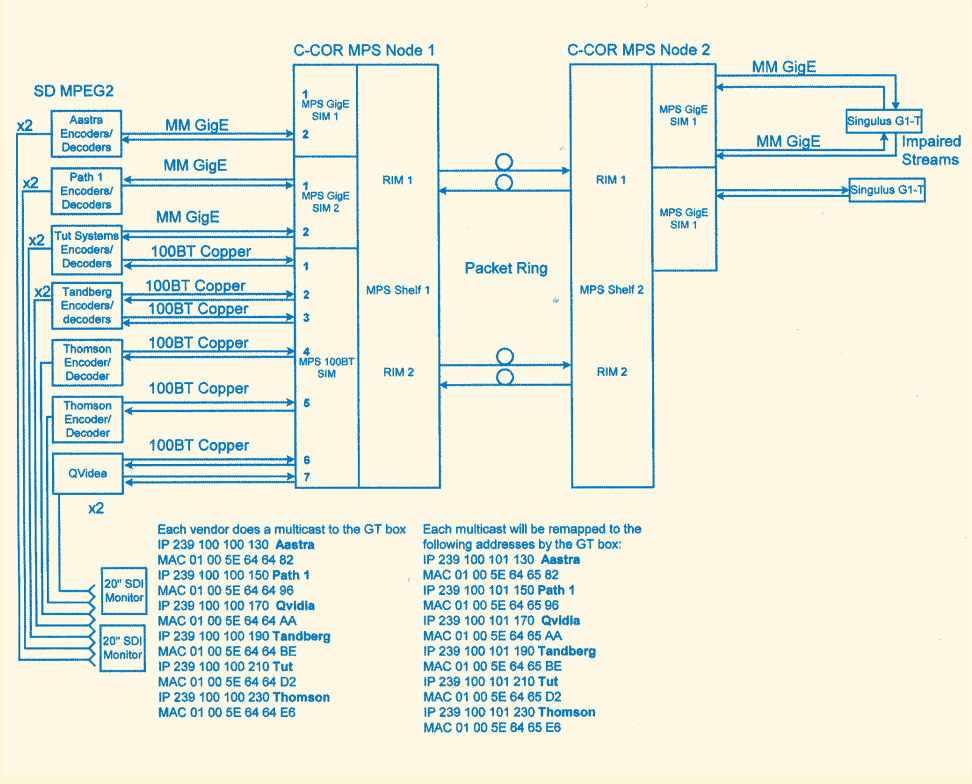

The following companies participated in the demonstration: Aastra Digital Video, C-COR, Path 1, IneoQuest Technologies, Inc., Qvidia, Tandberg Television, Thomson, and Tut systems. The figure above shows the demonstration equipment configuration.

The demonstration specifications were:

- 20.0 Mbits/sec MPEG-2 transport stream rate (CBR)
- 7 transport stream packets per IP packet
- Encapsulation as specified in Pro-MPEG Forum Code of Practice 3 Release 2
- 4:2:2 video with stereo MPEG audio
- FEC settings of Period = 5 and Order = 10 (correct 5 lost packets in a burst)

During the demonstration, significant interoperability between vendors was achieved. That said, this event was a trade show demonstration, not a lab interoperability test. As such, the demonstration results are not conclusive. Unfortunately, disruptions occurred during the demonstration period including switching and cabling problems, which prohibited the team from performing a thorough test. While disruptions are not uncommon in a trade show environment, conditions were not ideal for conclusive and detailed measurements.

In conclusion, the demonstration showed a high level of interoperability in a multivendor environment using the Pro-MPEG Forum Code of Practice 3, Release 2. Furthermore, the demonstration allowed conference attendees to see the difference between FEC and non-FEC on an impaired network. The VSF board is now looking into conducting a full lab test with a detailed test plan and lab style procedures in the future.

The next Video Services Forum meeting series will be held May 23-25, 2005 in Austin, TX. For additional information, visit the VSF website at www.videoservicesforum.org/Meetings/index.htm.