Distributed Encoding Architectures

By Kenelm Deen

The adoption of internet protocol (IP) technology into the broadcast industry has enabled a move from co-located encoding and multiplexing solutions to new solutions in which the encoding stage is pushed to the content provider. This removes the requirements for contribution encoding and re-encoding for final distribution. In a distributed encode solution, content is encoded with the philosophy "encoded once—distribute many." This saves an additional encode stage, conserves transmission bandwidth between the content provider and headends, and reduces network demand, lowering distribution cost while increasing capacity. Direct-to-home (DTH) headends typically operate statistical multiplexing (SM) to maximize quality and content delivery to the subscriber over the distribution pipe.

Centralized Architectures

Central headends usually consist of four distinct stages:

1. Content acquisition
2. Encoding
3. Multiplexing
4. Modulation

Content acquisition and modulation often occur at locations distant to the encoding and multiplexing and are typically connected by Telco networks. The encoding and multiplexing stages are traditionally co-located. In direct-to-home (DTH) applications, encoding efficiency is important—the lower the bit rate per service, the more services that can be delivered for the same cost over the distribution pipe. This drives DTH headends to operate statistical multiplexing (SM). Video content can be encoded using either constant bit rate (CBR) or variable bit rate (VBR). SM is the process of controlling VBR encoding to squeeze video components into a group bit rate. Bit rate is dynamically allocated across all components in the group, controlled by an SM controller. **Figure 1** shows the four stages typically found in a central headend.

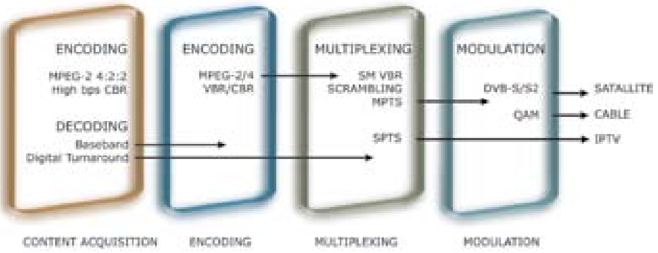

Figure 1. Stage within a central headend.

Content Acquisition

Content acquisition is the process of collecting content for encoding and delivery to the subscriber; it means different things to different people. It could mean collecting content off-air using banks of receivers for remultiplexing/re-encoding or could take the form of a contribution feed from a remote encode site feeding high-quality high-bit-rate content directly to the central headend for re-encoding. Of particular interest is contribution acquisition, in which the headend operates and controls the remote encoders.

Contribution feeds are mainly single program transport streams (SPTS), encoded using CBR. The feeds are encoded at a higher bit rate than the final distribution bit rate to minimize loss of quality in the second encode pass. Alternatively, multiprogram transport streams (MPTS) can be combined with SM to reduce distribution costs to the central headend (CHE). MPTS provides greater efficiency than SPTS, but adds complexity at the remote site. **Figure 2** shows a mixed headend architecture.

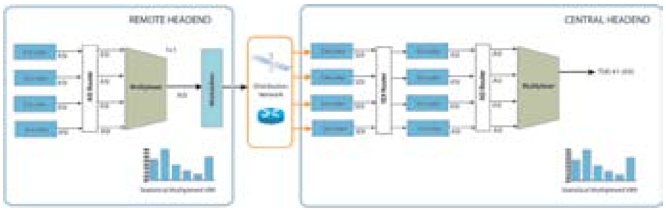

Figure 2. Mixed headend architecture.

Encoding Stage

The encoding stage encodes the content for delivery to the subscriber population. If the content was previously encoded, the bit rate or format may not be suitable for delivery to the home. Contribution feeds are normally encoded at higher bit rates than final distribution rates to maintain quality for editing. SM is common in CHE, maintaining picture quality while increasing delivery efficiency and enabling additional content delivery without increasing the delivery bandwidth.

The use of asynchronous serial interface (ASI) between the encoder and multiplexer is ideally suited to MPEG transport and SM systems in which the SM controller (often inside the multiplexer) controls the VBR encoding rate of a group of encoders, based on collective demand. This assumes closed loop SM systems, which deliver increased efficiency over open loop solutions in which the multiplexer has no control over the encoder. Open loop SM has limited ability to alter the bit rate while maintaining quality.

SM places stringent requirements on the interaction between the encoder and multiplexer and SM controller. This paper examines the impact of using IP networks as an interconnection technology on SM and highlights how these issues are resolved to provide high-performance SM over IP networks.

Multiplexing Stage

The multiplexing stage packages the content into a service for delivery to the consumer.

Multiplexing is normally associated with satellite and cable platforms where MPTS is used to deliver content to the home via shared distribution pipes on a broadcast model. In IPTV headends, content is distributed using CBR, or capped VBR SPTS because of restricted last-mile bandwidth pipes to the home.

The multiplexer stage will often incorporate security in the application of DVB-CSA (common scrambling algorithm) to the content for secure delivery to the home. The multiplexing stage generates a MPEG-2 transport stream for modulation and delivery.

Modulation Stage

The modulation stage applies the relevant modulation (DVB-S/S2 or QAM) for distribution to the consumer, and the modulated signal is then provided to the distribution system for delivery. The modulation stage is normally located with the distribution system and not co-located with the multiplexers.

ASI Transport

Asynchronous serial interface is the default standard for interconnecting broadcast compression equipment and is widely adopted within the industry. It provides the following characteristics for the carriage of MPEG compressed content between devices:

- Low (almost zero) latency.
- Deterministic bandwidth.
- Constant error rates.

These are ideal characteristics for compressed content, where bandwidth capacity is deterministic and there are known constant propagation delay and minimal error rates that are correctable using Reed Solomon error correction techniques. This adds minimal latency to the transmission and is transparent in terms of overall latency. This is not surprising, because ASI was developed for carriage of MPEG-2 transport stream packets for interconnecting broadcast equipment. ASI was the evolution of the original low-voltage differential signaling/serial parallel interface (LVDS/SPI).

Transport over ASI provides sufficient bandwidth for compressed content distribution with a maximum bit rate of 216 Mbits/sec. Combined with a low error rate correctable with Reed Solomon error correction with almost zero prorogation latency, it is ideally suited to transport MPEG-compressed content. The main restriction in using ASI to transport MPEG content is the 200-m interconnect limitation and expensive routing equipment. These restrictions affect today's business models, which have evolved requiring flexibility over static solutions that were the norm only a few years ago. ASI imposes restrictions on the design by high equipment cost and the fundamental requirement of co-locating devices within 200 m of each other.

IP Enters Central Headends

The industry is currently shifting from ASI to IP-enabled headends. Within the last 12 months, a shift from standard Telco networks such as ATM and DS3 to IP has occurred in the contribution market. This has fueled the acceptance of IP within the broadcast industry, and now a move to IP-enabled headends as begun.

The switch to IP-enabled headends is being driven by the increased flexibility in using IP over ASI and the reduction in equipment cost. The use of ASI routers and distribution amplifiers adds significant cost to headend designs while imposing interconnection and switching restrictions. IP inherently provides flexible switching and routing of content. The cost of IP network equipment is constantly falling, compared to ASI equipment, which is in a niche market. IP provides the following benefits:

- Increased capacity GigE and 10GigE over 216 Mbit/sec limited ASI capacity.
- No distance limitations.
- Full duplex (bi-directional) transport.

- Flexible routing.
- Improved redundancy solutions.

Compression vendors now include IP interfaces on their equipment, in addition to ASI, with the emphasis on IP connections moving forward. IP headends remove single points of failure, typically introduced by the placement of an ASI router between encoders and multiplexers for redundancy switching. This is done by enabling the multiplexer to select the content streams dynamically from any encoder on the IP network using standard multicast IP protocols.

This flexible content selection allows a shift in paradigm from isolated encoding pools dedicated to a single transport stream/multiplexer to one of a single encoder stage generating content. **Figure 3** depicts an IP headend containing equipment pools.

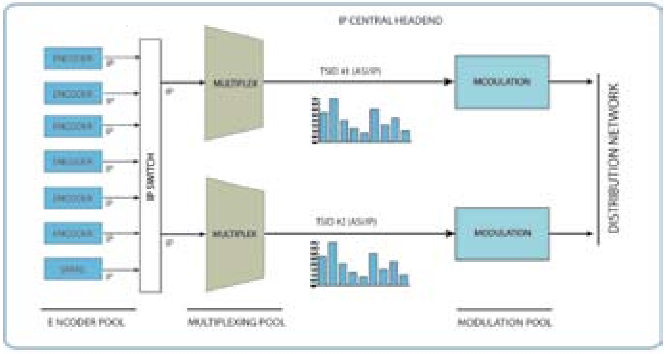

Figure 3. IP central headend.

HD the Driving Force

Broadcast high-definition (HD) content requires substantially greater transmission bandwidth for contribution feeds than standard-definition (SD). Transmitting HD uncompressed content requires 1.44 Gbits/sec compared to 270 Mbits/sec for SD, while the use of light compression standards such as JPEG 2000 reduce the bit rate to around 100 Mbits/sec, offering a substantial reduction of 1.44 Gbits/sec uncompressed 1080i HD, 100 Mbit/sec JPEG 2000 compressed HD, and 80 to 85 Mbit/sec MPEG-2 4:2:2 HD.

Contribution rate MPEG-2 SD requires approximately 18 Mbits/sec using 4:2:2 encoding to prevent chroma loss on re-encode. Although MPEG-4 AVC is the chosen format for delivery to the home, due to the low bit rate, there are currently no implementations of the MPEG-4 contribution profiles. This leaves MPEG-2 HD for contribution encoding, which requires 80 to 85 Mbits/sec for high-quality 4:2:2 transmissions. This high bit rate increases the network transmission costs, making the solution unacceptable for contribution/distribution of HD content between remote sites and the CHE.

HD broadcasts are nearly always VBR encoded to squeeze the maximum number of services into the available bandwidth. Typically, 3 or 4 HD services are broadcast in 38 Mbits/sec averaging 6 Mbits/sec per service in Europe and Asia and 4 or 5 HD services in the same bandwidth in the U.S. Taking this example, a contribution link would require a bandwidth of 340 Mbits/sec (4 x 85 Mbits/sec). This would require a STM-4 (622 Mbits/sec) pipe, which would be extremely costly and not easily accessible. This has forced broadcasters to place remote compression headends closer to the content, to minimize the contribution link bandwidth. This adds complexity and additional equipment costs and compromises the video encoding efficiency and quality as a result of encoding the content twice at similar bit rates. This is not ideal, but it is currently the only solution available.

Distributed Encoding to the Rescue

The use of distributed encoding over IP networks enables broadcasters to encode HD content if it is in the correct format (MPEG-4 AVC) for transmission to the CHE for multiplexing and distribution to the subscriber.

Distributed encoding benefits include a single encode stage, lower contribution bandwidth between sites, and simplified remote sites. This achieves the goal of maintaining the encoding efficiency, simplifying the equipment at the remote site, and reducing network costs.

Distributed Architectures

IP-enabled headends remove the requirement to co-locate the encoding and multiplexing stage, enabling the encoders to be distributed. This allows a broadcaster to place the final emission encoders at the content provider, removing the need to encode multiple times and reducing the transmission bandwidth required between sites.

In pushing the encoding stage out to the content providers, the need for traditional contribution encoding disappears. Now, only a single encoding stage is required, removing the need to encode twice while improving the quality and reducing the contribution bandwidth. A distributed architecture is shown in **Fig. 4**.

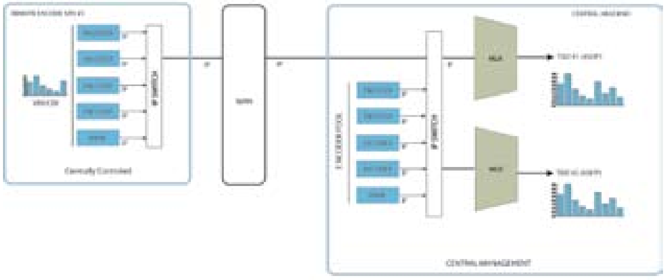

Figure 4. Distributed encoding architecture.

Instead of duplicating the functionality of the CHE at the remote site, the CHE becomes distributed, leaving only the multiplexing and possibly modulation. An IP network is used to link the encoding and multiplexing stage. In distributing the encoders, the IP network characteristics become more important. The move from a local area network (LAN) with specific performance characteristics to a wide area network (WAN) with variable characteristics emphasizes the effects of carriage over the IP network. Network-introduced impairments include high network latency (> 50 ms), packet loss and reordering, and variable inter arrival packet time (jitter).

These network characteristics add additional complexity to IP multiplexers because they must compensate for impairments; otherwise, service disruption will occur. IP networks can provide the required characteristics for transporting MPEG-compressed content but require careful design.

IP Network Characteristics

The carriage of realtime content places stringent requirements on the underlying network. A realtime broadcast requires constant bandwidth, consistent latency, and a low known bit error rate (BER). Traditional transport mediums such as ASI and ATM provide deterministic characteristics that are ideal for broadcast transmissions, whereas the Ethernet network is inherently indeterministic. The comparison is made with ATM technology because this was the preferred Telco network for contribution/distribution before the move to IP.

IP is a packet-switched technology, unlike ASI, which contained no framing and purely carries content between two points (point-to-point). In contrast, ATM is a circuit-switched technology that requires the setup of dedicated routes (connections between two nodes) for the duration of the transmission. IP packets are routed to their destination determined by network routing algorithms and do not require predefined routes to be established.

Transmission characteristics required for realtime content include 100% network availability, guaranteed bandwidth, consistent network latency, and defined error rate, allowing realtime recovery. If any of these are not met, service disruption will occur. For realtime transmissions, fixed bandwidth is essential for the duration of the transmission. If the network is unable to sustain bandwidth demand, service disruption will occur as packets are discarded. With network-introduced packet loss or variable latency, the receiving device will not be able to replace missing or delayed packets without affecting the service.

ASI was specifically designed for carriage of MPEG-2 transport packets, unlike IP networks, which were designed for general-purpose traffic utilizing a “best effort” approach. IP was defined for flexible routing; this is a key strength in its adoption within the broadcast industry, along with its flexible bandwidth allocations.

Broadcast IP networks are designed for the distribution of broadcast content with specific characteristics and sustained bandwidth availability. With the demand on IP networks to provide a consistent quality of service (QOS) comparable to that of traditional networks, a number of protocols can be applied to provide the required QoS.

Transmission Impairment

The encoder generates an IP packet flow that is routed over the IP network until it reaches the multiplexer. The process of transmitting the IP packet across the network may introduce impairment when the MPEG-2 stream is reconstructed. The following impairments are typical.

- Packet loss due to random or burst errors.
- Jitter from queuing within the network.
- Sequence errors (out of order packets).
- Variable latency between TX and RX sites.

The MPEG-2 transport stream (TS) packet is the preferred container for carriage of MPEG compressed content between broadcast devices. An alternative is the encapsulation of the elementary streams (ES) containing video, audio, etc., directly into IP packets, for delivery to consumer devices such as set top box (STB) and handheld devices.

The MPEG-2 TS packet must be encapsulated in an IP frame for transmission over the network. A standard IPv4 frame is 1500 bytes and an MPEG-2 TS packet is 188 bytes. This allows 7 MPEG-2 TS packets to be encapsulated into a single IPv4 frame. IPv6 frames are generally not used unless it can be guaranteed that the entire network is IPv6 enabled; otherwise, the larger IPv6 frames will be fragmented, resulting in jitter on the MPEG stream. RFC2250 defines the use of realtime protocol (RTP) for the transmission of MPEG content, because this allows ordering errors to be detected and possibly corrected.

Unless corrected, impairments result in service disruption. Depending on the level of impairment, correction techniques can be applied to recover and provide a stable transmission. One such technique is the application of forward error correction (FEC), which can allow the receiving device to recover missing packets. In addition, network protocols such as traffic classification and prioritization can be used to apply edge-to-edge QoS, based on traffic engineering (TE), to ensure bandwidth and minimize jitter and sequence errors.

MPEG-2 Transport Stream Recovery

The receiving device must recover the original TS transmitted by the encoder without the need for retransmission. The impairments introduced during transmission over the IP network will now be examined.

Packet Jitter

The received MPEG-2 TS will have jitter introduced during the transmission over the IP network. MPEG-2 TS packets are transmitted in IP frames that contain a maximum of 7 MPEG-2 TS packets.

This encapsulation process introduces a small amount of jitter, but carriage over the network introduces the largest jitter component. This IP jitter (inter packet arrival-time variance) results in MPEG PCR jitter, which, if not removed, may affect the service, depending on the level of jitter introduced. MPEG specifies the input signal should be within ±500 ns of PCR jitter and multiplexers normally are designed to remove a few milliseconds of jitter from the received MPEG-2 TS.

IP packets are not transmitted at a constant rate. Packets are sent in bursts between network devices as they transverse the network. If the network is experiencing high usage, it will hold packets in buffers, adding variable latency and resulting in

PCR jitter. The multiplexer must smooth out the introduced jitter. This is typically achieved using a de-jitter buffer, which absorbs input jitter by buffering the received IP packet flow and emptying the buffer at a constant rate equal to the original transmission rate. **Figure 5** depicts an IP de-jitter buffer and the smoothing process.

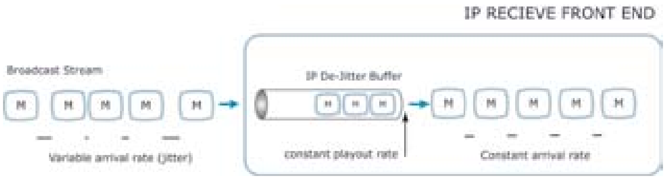

Figure 5. De-Jitter buffer model.

This requires the multiplexer to lock to the incoming input flow and determine the incoming rate in order to set the correct empty rate to remove the jitter. A reliable method of determining the TX rate of the IP flow is to use the 27 MHz PCR clock contained within the MPEG-2 TS packets. By taking successive samples from the PCR value, the output rate can be determined quickly and accurately. Alternative approaches derive the TX rate using the RTP 90-KHz clock, which has a lower resolution than the 27-MHz PCR clock, significantly increasing the locking time.

The size of buffer is dependent on the amount of jitter that must be removed. The addition of de-jitter buffers adds additional latency to the transmission path, which may be undesirable. Minimizing network-introduced jitter to reduce end-to-end latency through the transmission system is essential.

Latency

Latency is added between the encoders and multiplexer because of the propagation time of the IP flow over the network. This additional latency component forms part of the overall end-to-end system latency. This latency component cannot be removed once introduced. It is important to minimize network latency to keep end-to-end delay latency down and minimize redundancy switching time between remotely located encoders.

In CBR encoding, the addition of latency between the encoder and multiplexer does not add complexity to the solution. However, in a VBR SM solution, the design must compensate for the presence of this latency, otherwise the performance of the SM will be impaired depending on the amount of latency introduced.

Packet Loss

During transmission over the IP network, packet loss may occur. Missing IP packets will result in lost MPEG-2 TS packets and service disruption. Remember that a single IP packet can contain 7 MPEG-2 TS packets. A single lost IP packet causes a service outage. For this reason, the application of error correction data is essential for the recovery of missing packets.

The Pro-MPEG code of practice (CoP) FEC/SMPTE 2022-1 scheme defines an application in which additional FEC data streams are transmitted in separate IP flows to the content, enabling realtime recovery at receive devices. **Figure 6** shows the process of FEC transmission and recovery.

Figure 6. FEC transmission model.

The Pro-MPEG CoP FEC provides a flexible implementation, by which the level of error protection can be tailored to meet the network's performance. **Figure 7** shows the parameter trades possible. FEC provides burst and random packet loss recovery and variable protection levels. The use of FEC introduces additional considerations such as increased TX bandwidth and additional transmission latency due to the decode process.

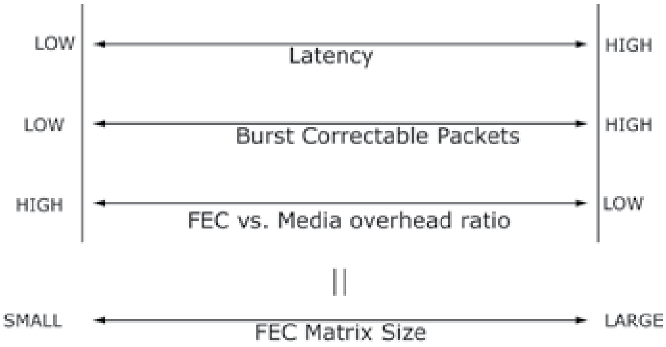

Figure 7. FEC parameter tradeoffs.

In CBR headend implementations, the addition of FEC between the encoder and multiplexer adds no additional complexity other than FEC processing. In VBR SM solutions, the addition of latency between the encoder and multiplexer has implications on the design of the SM system to handle the additional delay introduced by the FEC scheme. The effects of latency on the statistical multiplexer will be examined later.

Sequence Errors

Another major error common in IP transmission is the possibility of receiving packets in a different sequence than they were transmitted. This is termed sequence errors and results when the network sends different IP packets over different routes through the network.

RFC2250 defines the use of realtime protocol (RTP) for transmission, which includes a sequence number in the IP header, allowing a receiving device to determine out-of-sequence packets. **Figure 8** shows received sequence errors. In combination with the de-jitter buffer, this allows the receiver to reorder out-of-sequence packets. Excessive out-of-order packets cannot be recovered without the use of FEC, because the receiver can wait only a limited time before signaling a packet as missing.

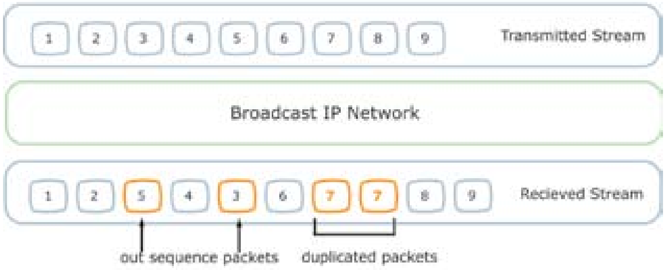

Figure 8. Receive packet sequence errors.

Statistical Multiplexing

SM is the process of VBR encoding multiple video components within a multiplex, allocating dynamic bit rate per a video component from a shared bit rate pool. There are two fundamental concepts—open and closed loop systems. In open loop, the downstream device has no interaction with the encoder and cannot adjust the encoded bit rate. In closed loop systems, the multiplexer interacts with the encoders and controls the encoders bit rate, based on encoder bit rate requests across all the encoders.

Closed loop SM will be examined, because it is the most common format for encoder multiplexer solutions. SM is essential because it enables a broadcaster to deliver additional services within the same transport stream bandwidth than if using CBR encoding. This reduces the cost per service by about 30%.

ASI-Based Statistical Multiplexing

SM has been around for many years. The first-generation solutions operated on proprietary interfaces between the encoder and multiplexer, such as transparent asynchronous transceiver interface (TAXI), because of the high-speed bi-directional properties by which data and control could share the same interface.

SM traditionally requires the constant low-latency MPEG interconnect and reliable high-speed, as well as low-latency communications channel between the encoder and multiplexer/SM controller. These characteristics are essential in changing the encoder bit rate within short accurate time intervals, to allow precise and co-timed changes in bit rate between various encoders.

With the introduction of ASI interconnects between the encoders and multiplexer, vendors moved to split solutions in which MPEG-2 TS packets were sent over ASI and SM messages over another interface, such as IP/optical/high-speed RS232.

Statistical Multiplexing Message Cycle

The following cycle takes place between the encoders and the SM controller. In an SM system, the encoders will encode using VBR at the bit rate allocated by the SM controller.

- The encoders will typically request an encode bit rate dependent on the difficulty of the source material based on an initial first-pass encode.
- The SM controller will allocate a bit rate to each encoder, taking the overall bandwidth available and the demands from the individual encoders into account. Some priority rules, such as target encode quality and bit rate limits will apply in setting the bit rate. The priority rules are often set dependent on service importance to the broadcaster.
- The encoder will then adjust the compression level to meet the assigned target bit rate set by the SM controller.

Figure 9 shows this continuous cycle. The interval between encoder bit rate request is driven by the need to alter the bit rate to ensure consistent encode quality, while efficiently allocating bit rate across all encoders within the group. This cycle of messages must happen quickly and reliably in order to maintain rapid bit rate changes.

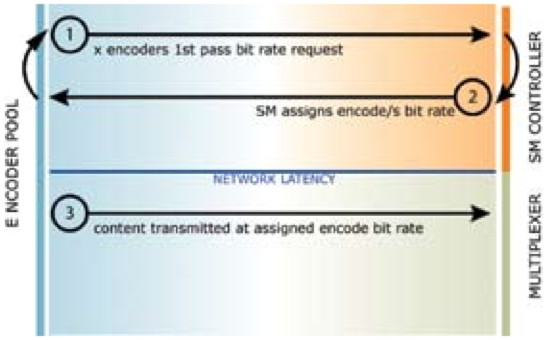

Figure 9. SM message cycle.

If it is assumed that the bit rate should change at the field/frame rate of the source content in order to maintain encode quality, taking an NTSC frame rate of 30 frames/sec would require changing the encoder bit rate every 33.33 ms, based on frame-level refresh, or every 16.66 ms for field-level refresh. This means that the roundtrip time between the encoders and SM controller must be less than 16.66 ms for message exchange and internal processing time and the MPEG-2 TS interconnects must have a latency less than 16.66 ms for the encoded content to arrive at the multiplex in time. It is essential that the interconnection between the individual encoders and multiplexer have the same latency, to ensure that the multiplexer receives the correct bit rate from each encoder at the correct time.

Figure 10 shows SM with equal latency paths. With ASI, the latency is almost zero and constant between all encoders. This assumes that all encoders are located in the same physical location, ensuring propagation delay is not relevant. These characteristics are easily achievable in headends with co-located encoders and multiplexers. This becomes an issue in distributed solution where the encoders are located in spilt locations and separated from the multiplexer.

LAN IP-Based Statistical Multiplexing

The introduction of IP networks connecting encoders and multiplexer within a CHE places stringent requirements on

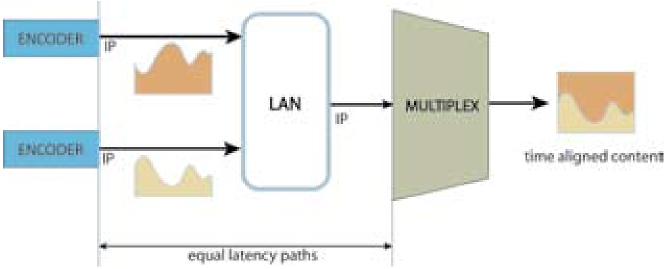

Figure 10. Equal latency paths.

the network's performance for SM. For SM, the IP network must have low-latency, unidirectional transmission from the encoders to the multiplexer and a bi-directional communication channel for SM messaging with a response time of less than 16.66 ms, in order to avoid imposing a performance restriction on the bit rate refresh rate. **Figure 11** shows different latency components in a CHE.

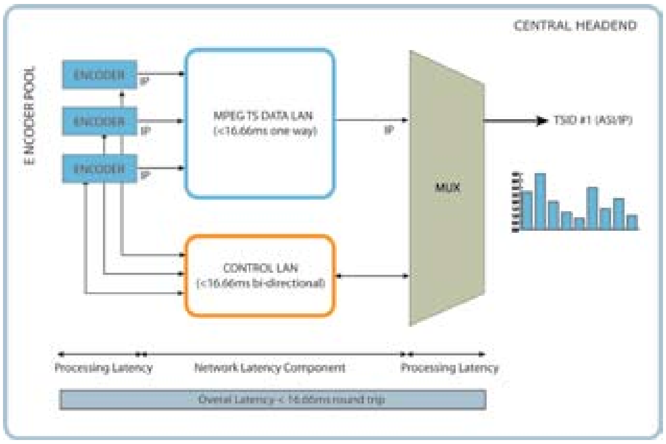

Figure 11. SM CHE timing.

Additional characteristics required for realtime content transmission include 100% network availability, guaranteed bandwidth, consistent latency, and defined error rate allowing realtime recovery.

Well-engineered LANs can easily meet these requirements, ensuring the same performance level associated with ASI implementations. GigE networks are normally used because they provide sufficient bandwidth without becoming congested, which would compromise performance. The LAN will be an isolated network solely for the interconnection of the compression equipment. A dedicated network should not introduce packet loss if correctly designed and does not require the use of FEC protection schemes. Based on the information supplied here, it is evident that IP-enabled headends can meet the performance demands required for SM.

WAN IP-Based Statistical Multiplexing

Operating SM over a WAN introduces additional complexity. The introduction of a WAN between the encoders and multiplexer can introduce latency in excess of 16 ms, variable latency paths between sites, and increased susceptibility to packet loss.

The introduction of increased latency exceeding 16 ms on the communication path will affect the rate at which the bit rate can be altered on the encoders if the same model used in a LAN is implemented. In order to maintain performance and rate of change of encoding bit rate, a different solution is required.

For a high-performance SM solution, the following issues need to be resolved:

- Latency equalization between encoders sharing a common bit rate pool, ensuring all content arrives at the multiplexer is time aligned.
- The SM algorithm adjusting for variable latency between the encoders and multiplexer and encoder and SM controller.

Latency Equalization

In SM headends, the latency between all encoders sharing a common bit rate pool must be the same. If the encoders do not share a common latency, the summation of all encoded bit rates arriving at the multiplexer will either overflow or underflow the group-allocated bit rate. An underflow wastes bit rate and reduces the efficiency, whereas an overflow breaks the transmission pipe, potentially resulting in service disruption. The effect of different latency paths between two encoders and a multiplexer is shown in **Fig. 12**.

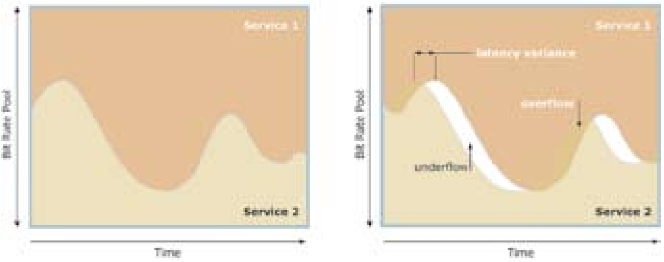

Figure 12. Latency path effects.

The effect of different latencies between encoders in the same bit rate pool and the need to equalize the latency to a common value can be seen. In co-located encoding and multiplexing headends, latency equalization is not required, because the latency is constant across all paths.

With the introduction of IP, headends containing local encoding and multiplexing do not require additional design changes because of constant latency across all encoders (assume same location). Most multiplexers are capable of removing a small amount of latency variance using the de-jitter buffer.

In a WAN implementation, if all encoders are located at the same location, it is likely that the delay will be reasonably consistent across all of the encoders. If the encoders are located in more than one location and are contained within the same SM pool, latency equalization is essential.

The latency variance between all encoders must be calculated and then removed to ensure content arrives time-aligned at the final multiplex engine within the multiplexer. Once the longest latency has been calculated, all encoders can be equalized to this value through additional buffering. It should be noted that this adds additional latency to the end-to-end system latency.

SM Algorithm

With the introduction of variable latency paths between the encoders and the multiplexer for transporting, the compressed content, and the increase in latency between the encoders and multiplexer for communication, the SM algorithm must be changed in order to maintain the SM performance associated with current ASI solutions.

Variable Latency Paths

The introduction of variable latency paths between the encoders and the multiplexer is the main complication with distributed SM solutions. The removal of latency is simple enough with the use of buffers to equalize the individual IP flows from the encoders. However, the calculation of the amount of latency between the encoders and the multiplexer is complicated.

A distributed SM solution must calculate the latency paths and allow adjustment to equalize the content delivery from the encoders at the multiplexing engine within the multiplexer. A common reference point is required between the encoders and the multiplexer in order to determine the latency of each RTP flow. The more accurate the calculation of the latency paths, the more accurate co-timed bit rate changes between the encoders can be managed.

The use of a common time source across the encoders and multiplexer provides a reference point for the calculation of latency paths. Once the latency paths have been calculated, all of the paths must be equalized to the longest latency path to ensure that all content arrives at the multiplexing engine co-timed. The calculation of latency paths can be static or dynamic, because these paths should not vary much over time. Only slight variance is expected because of jitter on the network.

Communication Latency Path

The latency between the encoders and SM controller is also important. The SM controller should always be co-located with the multiplexer or be aware of the worst latency path between the encoders and multiplexer.

The increase in latency between the encoders and SM controller must be accounted for in designing a solution. The encoders should constantly send bit rate requests to the SM controller and must have sufficient buffering to wait for a response on the encoding bit rate to use. The amount of buffering will depend on the latency between the encoders and the SM controller.

Conclusion

The introduction of IP networks into broadcast headends has not been without challenges, both commercially and technically. Broadcasters have had to gain confidence in MPEG transport over IP networks and are now embracing IP-enabled headends over traditional ASI infrastructures.

The main drive in adopting IP has been its ability to offer new headend architectures, which are more flexible than ASI solutions, removing single points of failure and simplifying the process of contribution encoding with the placement of distributed encoding. The introduction of distributed encoding has simplified remote sites by removing the need for two-stage encoding and multiplexing, reducing headend cost and lowering network transport bandwidth between facilities. This enables an “encode once–distribute many” model to be adopted.

IP networks are able to provide QOS service and availability required in the broadcast environment. The use of engineered IP networks combined with techniques such as SMPTE 2022-1 FEC allows reliable transmission over WANs.

The introduction of IP into headends has posed many technical challenges in providing the flexibility associated with IP, while maintaining traditional performance characteristics set by ASI. These technical challenges have been met, enabling broadcasters to take advantage of the benefits IP has to offer. The biggest challenge has been the introduction of latency and packet loss and the effect they have on SM and service delivery.

IP headends have arrived and are being deployed worldwide, replacing ASI solutions. The adoption of IP will eventually result in the disappearance of ASI as standard for broadcast equipment, but it is becoming an optional legacy interface.

Presented at the SMPTE 2007 Australia Conference in Sydney, Australia, July 17-20, 2007. Copyright © 2008 by SMPTE.

The Author

Kenelm Deen is a senior system architect within the Business Development Group of Tandberg Television, responsible for defining and validating architectures for satellite and terrestrial direct-to-home networks, contribution and distribution networks, and mobile TV for customer specific projects and future product development. Deen’s key areas of expertise extend across conditional access, headend architectures, IP, digital rights management, and MPEG-2/DVB standards. He has a MEng in electronics with computing and business from Kingston University in the U.K.