SOLUTION PAPER

Simplifying Headend Architectures

The All IP Approach

With the emergence of IPTV as a viable video entertainment distribution solution and the increased use of IP metro and wide area networks for video transport, more and more video processing equipment manufacturers are including IP interfaces on everything from encoders and receivers to splicers, rate shapers and test and monitoring systems. That, coupled with the existing wide array of high rate IP routers and switches, allows us to explore the advantages and mechanics of using IP as the primary connection paradigm for digital headends.

Advantages of IP

The advantages of using IP for connectivity within a headend include simplified wiring, lower implementation costs and enhanced redundancy control.

Simplified wiring - The addressing and routing capabilities of IP streams and the ubiquitous availability of gigabit Ethernet means that many compressed video/audio streams can be carried inexpensively over a single wire. In a headend, IP-ready video processing equipment can be terminated at an IP switch to aggregate compressed content feeds before transporting the signals. Today, using MPEG-4 which yields standard definition streams at about 2 Mbps and high definition at 6 – 10 Mbps, it's easy to see how an entire channel lineup can be transported over a single gigabit Ethernet connection.

Lower implementation cost – When all of the connectivity in a headend is based upon IP technology, less expensive IP switching and routing equipment can be used in place of ASI or SDI routing systems. In addition, IP signals can be transported reliably over longer distances than ASI or SDI signals thus eliminating the need for distribution amplifiers. Examples of these types of savings will appear throughout this review as the various IP-ready video processing elements are described.

Enhanced redundancy control – The connectionless packet switching capabilities of IP interfaces enable enhanced redundancy models previously unavailable with hardwired point to point (HD)-SDI and ASI interfaces. In a pure IP headend environment, all of the video processing equipment has access to all of the video content at all times without the need for any physical connection switching. This enables new redundancy paradigms where the redundancy architecture is no longer constrained by the physical signal flow.

Traditional Architecture

Consider the traditional redundant headend architecture. A traditional architecture relies on a variety of processing components connected via (HD)-SDI or ASI video routers. Here, all signals must be received and converted to (HD)-SDI baseband signals prior to being re-encoded in the desired format, such as MPEG-4 AVC.

A traditional facility will require a variety of ingest equipment such as satellite receivers, off-air receivers, and possibly fiber transport to ASI converters. These signals are all processed to a baseband signal and then distributed, with an audio/video switch, throughout the facility and to the encoding devices. The cost associated with this approach can be quite extensive when taking into consideration the complexities of designing, installing and maintaining such architectures. Additional expenses will be incurred in rack space and power required to operate a facility with all of this equipment.

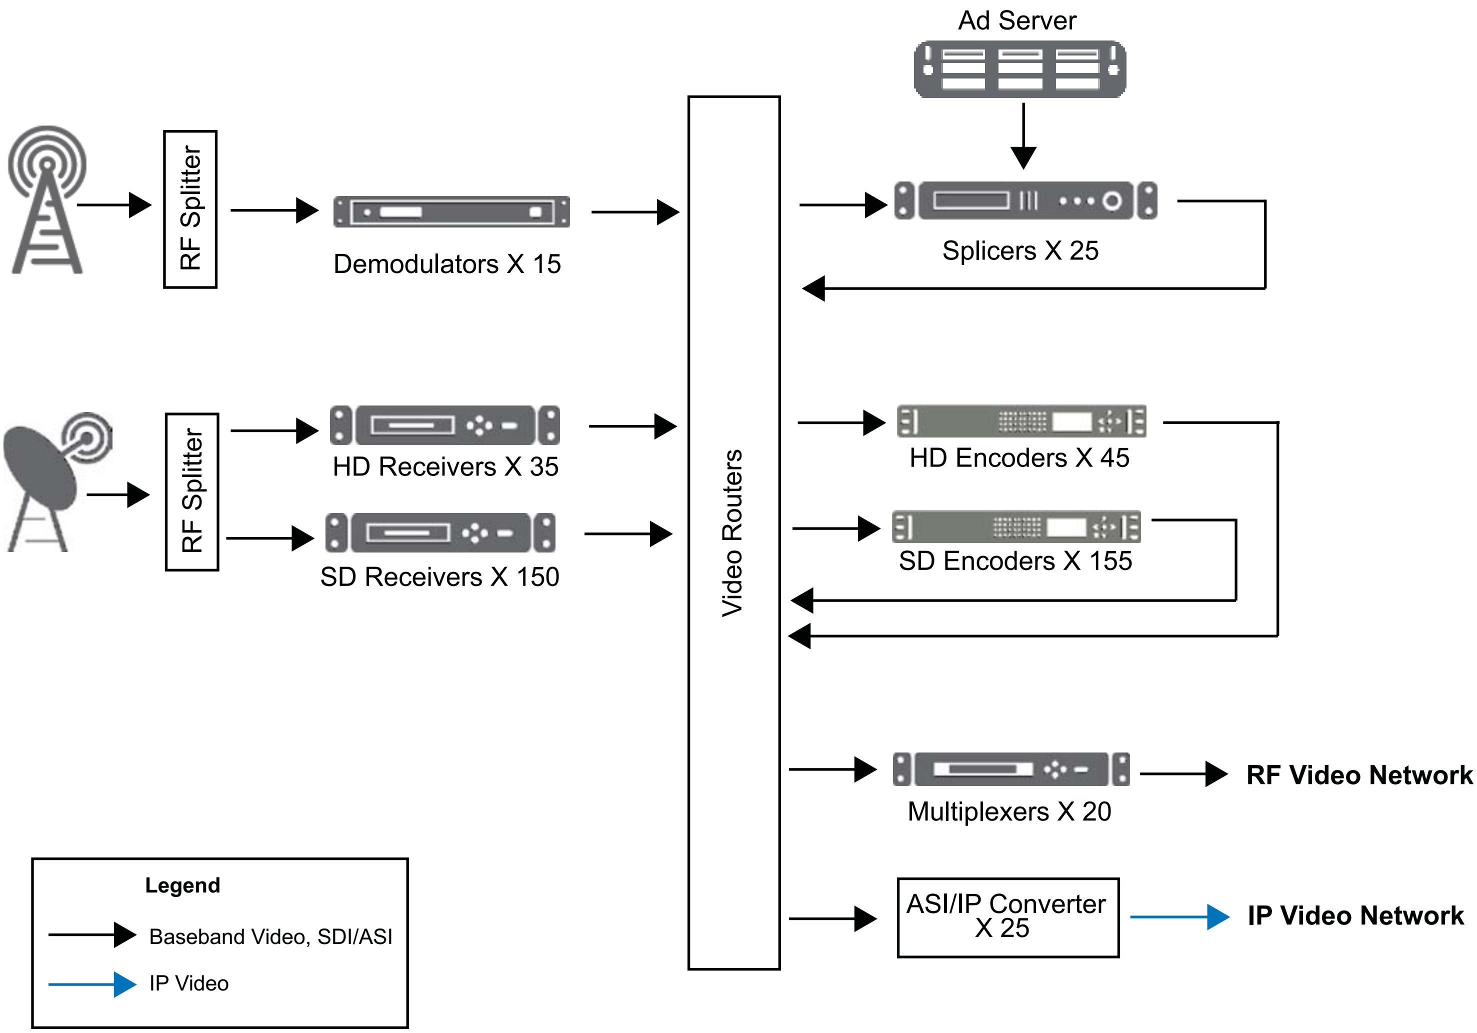

FIGURE 1: TRADITIONAL ARCHITECTURE

Figure 1 illustrates a two hundred channel redundant headend architecture. The channel line up consists of 35 high definition (HD) channels, 150 standard definition (SD) channels and 15 off-air ATSC/NTSC channels, 10 of which are HD and 5 SD. This configuration also includes targeted advertising for 25 of the 200 channels.

In a traditional redundant architecture, incoming video from satellite receivers and off-air demodulators is delivered to the appropriate encoders and splicers by a video router. The video router ensures that the correct video signals are presented to each device during normal operating conditions and in the event of a device failure. Once the splicing and encoding is complete the video must pass back through the video router before it can be aggregated for RF or IP video distribution. This traditional architecture using (HD)-SDI and/or ASI video transport and requiring video routers throughout the headend uses a total of 1115 cables, 470 devices and has 1870 potential points of failure.

IP Architecture

Comparing the simplified redundant IP headend to the traditional architecture the complexity is reduced by a factor of 4. As figure 2 illustrates, less equipment and cabling is needed than in the traditional architecture.

To further simplify headend video processing, the latest generation of headend equipment offers IP interfaces while consolidating multiple functions into a single device. Two examples of this consolidation are encoders that incorporate off-air ATSC/NTSC receivers and satellite receivers that incorporate transcoding capabilities. The off-air ATSC/NTSC receiver encoder eliminates the need for an external demodulator and the encoders IP output capability eliminates the need for a video router. Replacing the video router with an IP switch also eliminates ASI to IP converters, further reducing the total number of devices. Utilizing the latest generation satellite receivers with multi-channel functionality reduces the number of receivers as well as cabling. An obvious benefit is the reduced number of cables needed in a pure IP headend compared with a traditional headend. Our traditional example used over 1115 cables, while the IP architecture only uses 323. Fewer cables imply that there are fewer points of failure in this model.

FIGURE 2: IP ARCHITECTURE

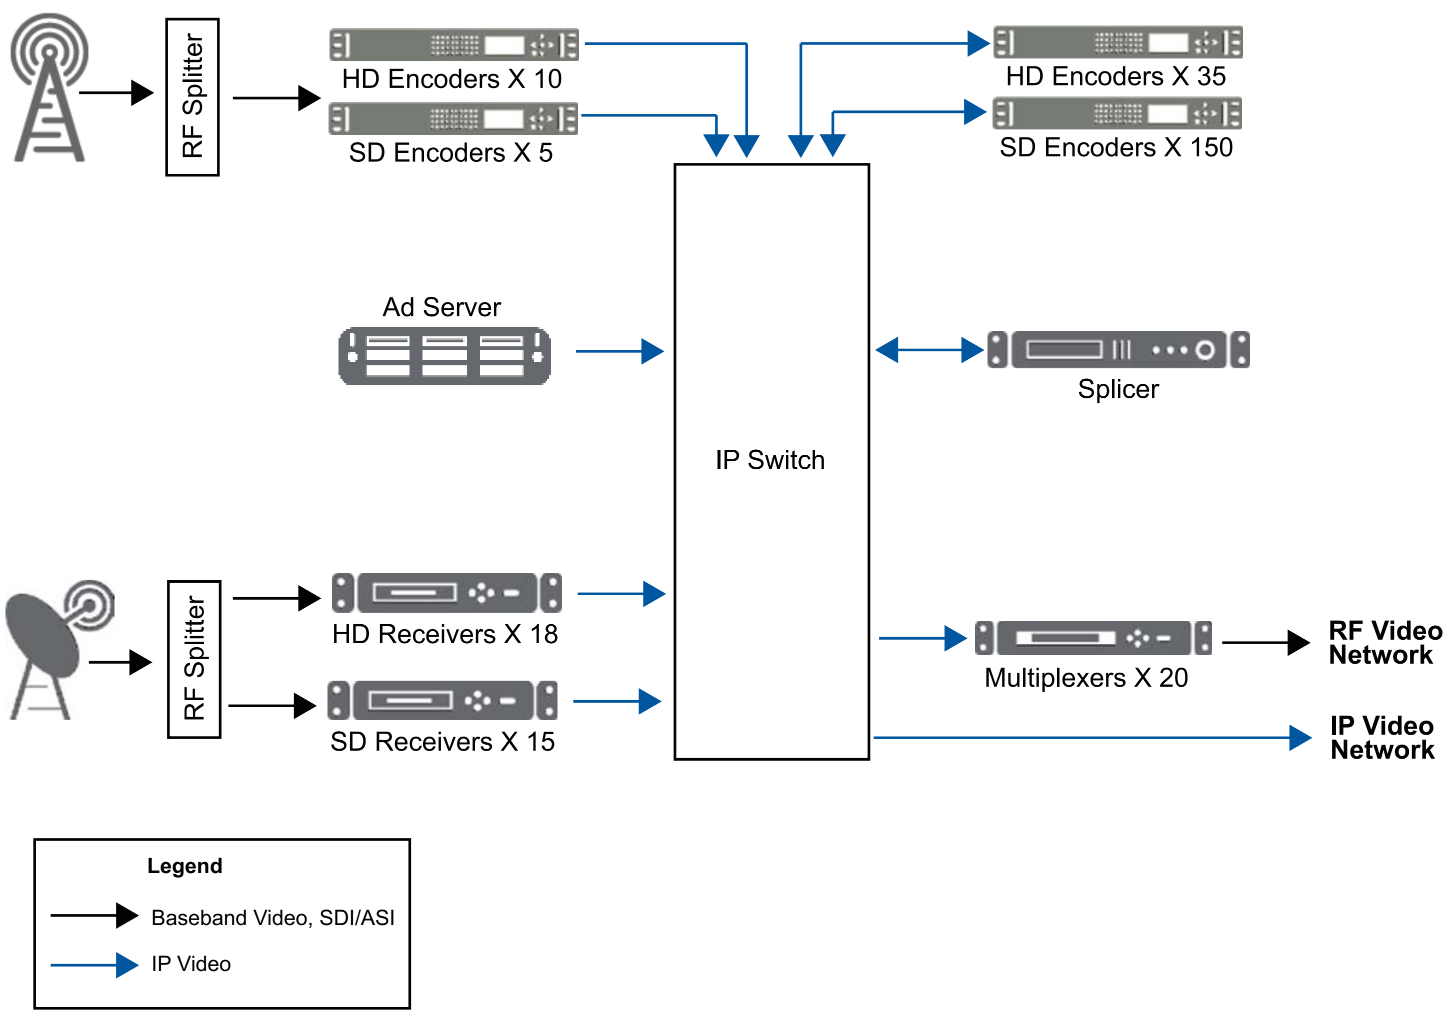

FIGURE 3: COMPARISON OF ARCHITECTURES

	Traditional	IP
Number of Cables (baseband and Ethernet)	1115	323
Potential Points of Failure	1830	203
Number of Devices	470	274

A simplified comparison between redundant traditional and redundant IP headend architectures is illustrated in Figure 3.

Using an all IP headend configuration significantly reduces the number of cables, devices and potential points of failure. Decreasing the number of devices lowers implementation cost, which lowers capital expenditures (CAPEX) and operating expenditures (OPEX). Deploying less equipment reduces power consumption, which has a positive impact on the environment. Simplified wiring, lower implementation costs and enhanced redundancy control are the main advantages of using IP connectivity in a headend environment.

Motorola can now deliver a completely IP based solution including satellite receivers such as the DSR-4410MD, grooming and splicing solutions such as the CAP-1000 and encoding solutions such as the SE-5150 Receiver/Encoder and the new NE-4112 Dual Channel Network Encoder. Maintaining an all IP headend infrastructure dramatically simplifies the implementation and significantly reduces cost. Traditional SDI and ASI video interfaces can only carry video in one direction. The IP interfaces are inherently bi-directional which allows the input video to be carried on the same physical cable as the output video. The advantage is that half as many cables are needed to carry video within the headend. Additionally, encoder redundancy which traditionally required separate ASI and SDI video routers - can now be implemented with the existing IP switch.

With all video being transported via IP in a compressed mode, health and status information in addition to video pictures can easily be sent over a facilities local area network to workstation-based monitoring and control sites. Motorola encoders offer a feature where a “thumbnail” representation of a high definition picture can be simultaneously output from the encoder on a separate IP address. The thumbnail can be routed to monitoring devices so that operators can see actual pictures in addition to stream and picture parameters.

IP Receivers

Motorola Commercial Integrated Receiver/Decoder (IRDs) and bulk descramblers have included digital outputs via ASI and now incorporate IP ports for monitoring, control and video transport. With these new ports, IRD outputs can be connected directly to an IP switch where the compressed streams can be routed to transcoders (for converting MPEG-2 to MPEG-4) or rate shapers (for reducing stream bitrates) or directly to the aggregation switch if the streams can simply be retransmitted in their native form and format. In the case where only legacy IRDs can be used, there are a number of ASI to IP converters that can inexpensively aggregate multiple ASI streams into a single gigabit Ethernet transport. In addition to the video, the health and status of the IRD can be reported via SNMP over IP and connected over the same LAN as the video streams or over a VLAN to minimize congestion.

IP Encoders

Motorola has implemented several features within the MPEG-4 encoding products which eliminate equipment and simplify headend architectures. The SE Series of IP capable MPEG-4 encoders incorporates an off-air ATSC/NTSC receiver eliminating the need for a two box solution. The NE Series MPEG-4 encoders take the IP approach a step further by delivering multi-channel video compression capabilities in a pure IP network device.

IP Splicers

The Motorola CAP-1000 provides several essential applications on a single platform including ad splicing, rate shaping, stat muxing and grooming. It has four available GigE IP interfaces, and manages a combination of up to 256 SD or 64 HD streams. The multi-channel form factor reduces the number of devices needed while providing a high density solution. Unlike traditional architectures where ad servers are connected directly to the splicers, the CAP-1000 connects directly to an IP switch, as does the ad server. This is beneficial because all ads stored on the server are available on the network at all times. Another function of the device is advanced overlay; utilizing IP, these capabilities occur simultaneous to splicing.

Applications

The following sections illustrate how incorporating receive functions into the MPEG-4 encoder and migrating to an all IP headend environment dramatically simplifies architectures and integrates functions.

Off-Air ATSC Receiver Encoder with IP Interface

Where local headends supply MPEG-4 formatted streams to the access network, there will always be the requirement to translate the supplied MPEG-2 broadcast feeds to MPEG-4 format. In this application, the system tunes to the appropriate frequency and program number to select the high definition (HD) or standard definition (SD) service. The receiver encoder demodulates the signal and presents the compressed MPEG-2 stream to the processing core. The video signal is decoded and re-encoded into MPEG-4 AVC. When necessary the HD signal is down-converted to SD following the active format descriptor (AFD) to maintain the proper aspect ratio. The audio components are optionally passed through or re-encoded into a new format.

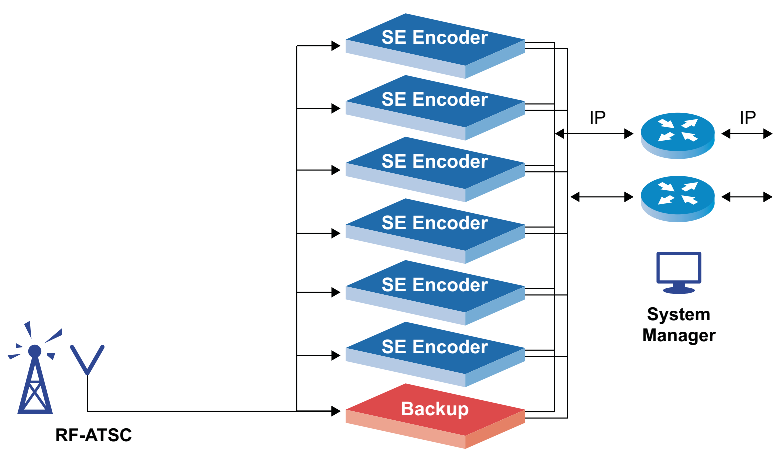

NTSC to AVC Receiver Encoders

Satellite Receiver with IP Interface

Direct reception of incoming satellite feeds requires a different approach than off-air NTSC/ATSC signals. In many places, proprietary solutions exist, which require the deployment of corresponding proprietary receiver/descramblers or, as is the case in DVB-S environments, common receivers utilizing internal proprietary descramblers.

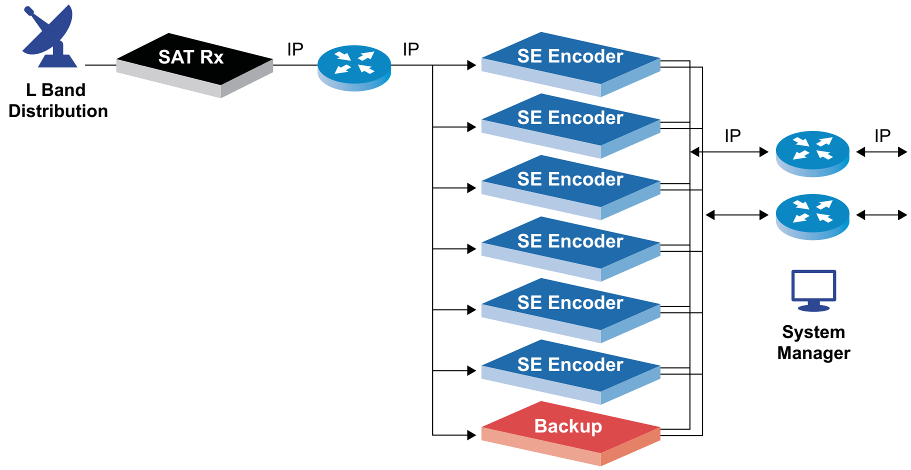

MPEG-2 to AVC Receiver Encoders

In these scenarios the MPEG-2 signal from a satellite transponder is decrypted and presented as a compressed MPEG-2 transport stream (TS) over IP. The encoder simply selects the desired service from the multi-program transport stream (MPTS) and delivers this compressed MPEG-2 transport stream into the AVC processing engine for decode and re-encode. Compressed audio is optionally passed through or re-encoded into a new format. The functionality includes the ability to deliver appropriate closed captioning and other peripheral signals.

Statistical Multiplexing over IP

Closed loop statistical multiplexing over IP allows the complexity analysis and rate allocation messaging to be delivered over the IP network. In a closed loop system each encoder reports the complexity of the video to a control application that then computes and instructs the encoder an appropriate bit rate to use. This technique enables the encoders to aggressively flex the bit rate to optimally track complexity and hence deliver better efficiency while maintaining video quality.

The IP-based architecture offers significant benefits compared to traditional statistical multiplexing solutions including reduced cost by eliminating expensive ASI infrastructure and replacing it with low cost IP equipment and the ability to allow encoders in multiple physical locations to participate in a single stat mux group.

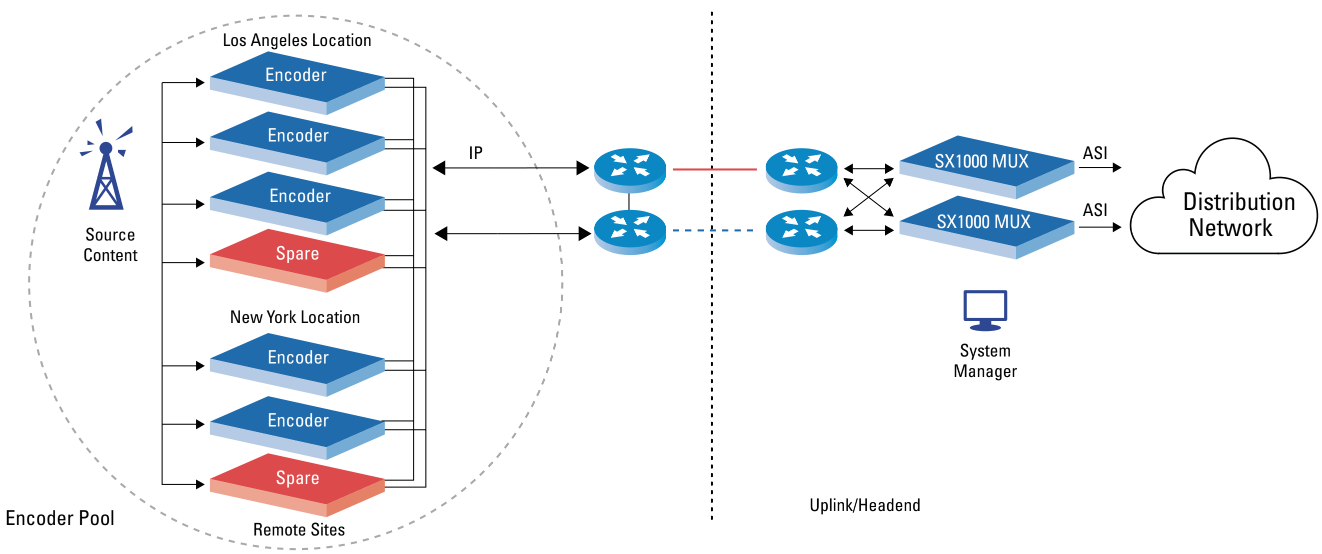

Conclusion

In the past few years, a wide array of video processing equipment has been outfitted with IP interfaces for monitoring and control so that element management systems could easily incorporate these devices within their system management umbrella. Now, as IP switches and routers that can handle gigabit transport have become widely available, these same video devices are adding the capability for the processed video to travel over IP networks as well.

This migration to ubiquitous use of IP for both control and transport has opened the door for the implementation of cost effective, compact headends that are easier to operate and maintain and less expensive to deploy.

MOTOROLA

www.motorola.com/ipvideo