IP VIDEO TRANSPORT SOLUTIONS FOR CABLE OPERATORS

Mark Davis
S. V. Vasudevan
David Brown
BigBand Networks, Inc.

Abstract

With content and service expansion a constant in the cable industry, operators have had to adapt their plants for increasing capacity demands, while also containing costs for maximum return on investment. This has driven several cycles of evolving transport techniques to deliver content from centralized sources to distributed subscriber bases.

IP technologies have extended new promise and seen initial implementations that achieve superior economics and functionality for data and voice, and increasingly video content. The openness and flexibility of IP enables rapid advances and cable operators can now contemplate the economic utilization of fiber distribution at both regional and national levels. There are several courses to consider for more widespread utilization of IP technologies for the distribution of core cable programming, and this paper considers their relative advantages and challenges, while recommending commitment to a highly flexible infrastructure able to adapt to emerging opportunities.

VIDEO TRUNKING EVOLUTION

Analog satellite was the first "trunking" technology used to cost effectively deliver multiple video programs over long distances from a single point to multiple destinations of nationwide headends. But evolutionary drivers proceeded quickly, as cable operators continued to see capacity requirements explode along with operator consolidation in the 1980s and 1990s.

Needs emerged to retool the metropolitan network architecture from standalone headends to headend-hub architectures. Most cable operators maintained wireless practices by leveraging CARS band microwave for video trunking between these facilities, until fiber optic transport became the new industry standard capable of transmitting analog video over relatively long distances without sacrificing picture quality. This optical transport method utilized FM and AM modulation techniques.

The AM super-trunk eventually emerged as the low-cost reliable solution and is still used for many hubs today. The advantage has been that signals could be processed or modulated at the central headend and then distributed to hubs up to 30 miles away over a single mode fiber. 1550nm optical transmission technology evolved with higher output lasers and EDFA amplifiers which pushed distances up to 90 miles. A simple low-cost optical to electrical receiver kept space requirements to a minimum and allowed retransmission to optical nodes in neighborhoods, commonly known as hybrid fiber/coax today.

However, system consolidation and the pressure to cut more operating costs and add more channels continued through the 1990s, which in turn drove a new digital transport technology. Systems emerged using proprietary encoders integrated with SONET-like transport, allowing delivery without the picture quality challenges or distance limitations of AM super-trunks. In addition, small decoders converted digital video streams back to IF frequencies and allowed the use of small low-cost IF-to-RF upconverters. Limitations included only 16 channels per fiber capacity and proprietary techniques that constrained interoperability and concentrated market power with particular vendors. SONET and even ATM vendors did make several attempts to solve these issues and were able to capture a small number of sites, but primarily failed due to the costly and heavy floor space requirements of baseband to RF conversion or modulation required at all hubs.

Thousands of SONET-oriented terminals were deployed at distant digital hubs. The industry then launched digital programming to the home in the late 1990s, requiring expensive digital video headend and video processing equipment. AM super-trunking remained a compelling solution through the transition to MPEG video because of its cost effectiveness. However, limitations on transport distances and functionality such as program grooming and hub-based local insertion still required some use of transport terminals until now.

IP TRANSPORT EMERGES

Recent widespread deployments of VOD (Video On-Demand) has driven new requirements for transporting large quantities of video streams and high volume QAM modulation at the hub or edge of the network. This has prompted reconsideration of the industry's video transport techniques. Early VOD deployments relied on costly ASI transport and distributed VOD servers. Technology vendors responded with new low-cost, high-density QAM modulators with integrated IP-to-MPEG-2 decoders that used Gigabit Ethernet as the transport layer. These edge QAMs were deployed using a variety of optical connections including DWDM (Dense Wavelength Division Multiplexing) and CWDM (Course Wavelength Division Multiplexing) technologies. New standards-based SFP (Small Form-Factor Pluggable) optical transceivers have allowed transport and edge QAM vendors to drastically cut the costs associated with both CWDM and now DWDM transport. In addition, GigE over next generation SONET as well as IP routing solutions costs also plummeted. These events opened up a host of new standards based, low-cost, and flexible video transport solutions.

Pushing digital programming deeper into the network has improved picture quality and reliability, lowered operating costs by consolidating headends, and provided capacity to meet ever growing programming requirements. Catalyzed by VOD requirements, utilizing IP technologies for this deep digital transport has opened new opportunities for the industry.

As digital broadcast channel capacity continues to increase, an operator has to consider the economics of transitioning away from traditional 1550nm super-trunks or continuing to expand proprietary digital transport. With the drastic decreases in costs outlined above, transitioning all digital programs to MPEG-2 over IP/GigE can be achieved at similar costs to proprietary digital system expansion, with significant benefits from these open and highly functional standards.

A more recent driver of digital transport evolution and bandwidth management requirements is the introduction of digital simulcast. This is the digital encoding of all analog channels, along with QAM modulation of the traditional analog tier that is then simultaneously, along with a decoded version of the analog tier, transmitted to subscribers. Encoding of all remaining satellite delivered analog programs along with hundreds of PEG (Public access, Educational and Government) channels is required for digital simulcast.

Simulcast allows operators to start the migration to an all-digital network and provides multiple near-term benefits such as use of less expensive, all-digital STBs (Set-Top Boxes), and improved picture and audio quality especially to high-end televisions. Simulcast also drives the transition of analog commercial ad insertion to an all-DPI (Digital Program Insertion) solution. DPI servers can be centralized and upgraded to provide GigE outputs, allowing DPI or digital splicing anywhere in the network, including hub-based ad zones. Simulcasting leverages the benefits of Gigabit Ethernet and IP transport techniques, advancing their predominant establishment in cable networks for all media and services.

Industry consolidation is continuing to drive the need to further consolidate and convert headends to hubs in larger metro areas and even fiber-connected rural areas. Many smaller headends in sparsely populated areas are now connected with leased or owned fiber for high speed data and telephony applications.

As the number of subscribers served by a headend increases, operators must consider adding a second redundant headend on the backbone in order to ensure the highest reliability or network availability, while delivering the lowest possible cost service. MPEG-2 over standards-based GigE and IP allows operators to utilize MPEG aware switching platforms to perform automatic failover or improve fault tolerance at the program level.

As transport trends progress, the metro concept is expanding and the benefits of nationwide terrestrial transport based on GigE, IP and optical technologies is becoming apparent. Several operators now have substantial wide-area fiber resources, or can cost-justify lambda leases from long haul fiber providers and consolidate all video, voice and data traffic onto a common backbone. Smaller operators are also faced with all the same challenges and have started forming partnerships to build a common channel IP video backbone fed by redundant primary headends.

FUNDAMENTALS OF MPEG-2 AND IP VIDEO NETWORKING

MPEG-2 transport as standardized by ISO 13818-1 has established itself as the de-facto protocol for the carriage of broadcast digital television services. Even the newer advanced video codecs such as H.264 can be carried over MPEG-2 transport. A coded video or audio frame is fragmented into several MPEG-2 transport packets, which are a fixed 188 bytes in length. A 13-bit PID (Packet Identifier), present in every MPEG-2 transport packet, identifies the elementary video and audio packet streams that comprise a broadcast television program. By using different PIDs to separately distinguish different elementary streams, an entire program can be multiplexed and carried as an SPTS (Single Program Transport Stream). Expanding this methodology, several programs, each with their own unique video and audio PIDs, can be broadcast as a unified multiplex, known as a MPTS (Multi-Program Transport Stream).

To carry MPEG-2 transport over IP networks, a further encapsulation step is required to place the MPEG-2 transport packets in an IP networking envelope. UDP (User Datagram Protocol) is most commonly used to carry both broadcast and narrowcast MPEG-2 traffic. This has advantages for real-time content like video over widely known TCP (Transmission Control Protocol) by eliminating the need for packet-receipt acknowledgements back to the transport source. Since the 24 byte UDP header represents additional transmission overhead, it is prudent to carry as many MPEG-2 transport packets in a single UDP message as possible. Typically, seven MPEG-2 packets are placed in a UDP message. This number is chosen because it represents the maximum UDP message size that fits into the maximum 1,500 byte payload size that is dictated by the 802.3 Ethernet framing format. While larger UDP message sizes are possible, limiting the overall payload to less than 1,500 bytes guarantees that the message will not be subjected to any unnecessary IP fragmentation procedures that could be asserted by lower-level protocol stacks.

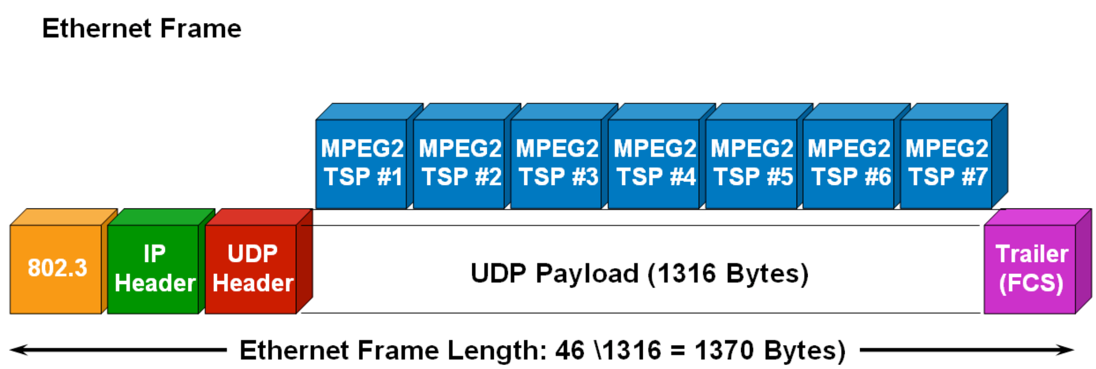

Figure 1. Encapsulating MPEG-2 content within Ethernet frames achieves very low overhead consumption and supports dozens of QAM multiplexes per link.

MPEG-2 is sensitive to jitter, which is variation in delay in the arrival of transport packets. This is because the transport stream carries timing information that is used by the receiver (e.g. an STB) to faithfully decode and regenerate the baseband program of interest. In an end-to-end digital television system, a 27 MHz system clock is typically locked to the incoming baseband video stream, although it is also possible to originate a stream with a free-running clock. This 27 MHz system clock drives a counter that generates a 42-bit PCR (Program Clock Reference). PCRs are inserted into the transport stream at a regular rate (at least 10 times/second), and are in essence real-time snapshots of the counter in the encoder. Decoders receive and extract PCRs for the particular program they are decoding and use the values and appropriate filtering to drive a VCXO (Voltage-Controlled Crystal Oscillator) on the receive side to accurately regenerate the 27Mhz clock, which is then used to regenerate the baseband video timing. This transmission model allows a decoder in a consumer's home to regenerate a video clock that is locked in phase to the encoder that originally compressed the program, which could be in a satellite uplink facility thousands of miles away.

Minimizing PCR jitter is important to maintaining a robust and high-quality end-to-end signal. ISO 13818-1 specifies a maximum PCR jitter of 500 ns, but this value specifically does not include jitter that can be caused by UDP/IP encapsulation and network transport by IP or other protocol. It is not uncommon to observe UDP/IP-encapsulated network traffic with tens of milliseconds of PCR jitter, almost two orders of magnitude above that maximum limit specified by MPEG.

To cope with this high network jitter, most video-aware networking components de-jitter the stream by intelligently correcting the PCR value in the transport packet and/or shaping the flow of the packet traffic as it transits through the device. Without an effective de-jittering mechanism, excessive PCR jitter can cause packet deliveries that violate the buffer models specified by MPEG, and more importantly can prevent the decoder from accurately regenerating the 27MHz clock required to reconstruct the baseband signal. Consumer televisions that encounter this type of impairment will generally be unable to lock to the corrupted "color burst" waveform that precedes the delivery of a field of NTSC video – the most common symptom of this situation is that the video picture will be unable to render any chrominance information and becomes "black-and-white".

Another approach to de-jittering is the use of RTP (Real-time Transport Protocol), as specified by IETF RFCs 1889 and 2250. RTP is a popular protocol for streaming media applications, and has mechanisms to support multi-source applications such as videoconferencing. A UDP message would be the data payload of an RTP packet. An RTP header contains a timestamp that can be used to recover and restore packet timing between a network sender and receiver. However, jitter that is induced by the UDP encapsulation of MPEG-2 transport packets cannot be recovered with RTP timestamps, and RTP adds another 12 bytes of overhead to the overall message size.

NEW VIDEO BACKBONE REQUIREMENTS

IP capabilities are becoming increasingly applicable for video traffic in cable networks at an opportune juncture of both service expansion opportunities and competitive threats to the industry. Customer tolerance to network outages will quickly diminish now that the competitive landscape is changing with telcos and satellite video providers vying for market share. Given this risk, any video network architecture design must have very high availability, fault detection and rapid recovery features.

Increasing channel counts along with more HD programming will continue to drive the need for additional capacity. Channel expansion must be easy to implement while minimizing costs.

Pressure to minimize headcount expenses while driving high network availability, will require networks to remain simple to provision, operate and easy to troubleshoot. Minimizing the number of appliances and network elements such as small-profile but limited-purpose "pizza boxes" minimizes the amount of training and device dependencies. Complex routed networks can also be difficult to troubleshoot when a number of multi-service traffic related challenges like denial of service attacks appear on the video network. Safeguards and sound QoS (Quality of Service) practices must be implemented to avoid these challenges.

Now that the billion dollar access network upgrade is complete, there is financial pressure on MSOs to drive free cash flow by limiting capital spending. This will drive efficient use of existing infrastructure like fiber and existing data transport elements. WDM technologies are an attractive way to extract maximum returns on existing specialized network element investments while allowing growth and graceful transition to an all IP video transport network. There is no need to do a fork-lift upgrade of all transport devices with WDM. Upgrades can happen when service specific requirements dictate as with the case of video transport expansion requirements.

The renowned flexibility of IP expands cable operator choice in how transport networks are configured and operated. A combination of business and technology drivers are positioning the industry to consider much more widespread transport implementations than the classic headend-hub metro connection with nationwide digital distribution. Several implementations of this are available for MSO consideration. Robust and flexible infrastructure proves key to maintaining best-of-breed options over time while implementing the best techniques to address current needs.

PASSIVE WDM NETWORKS

Passive WDM systems are now readily available for a variety of network transport systems. Passive WDM utilizes industry standard ITU (International Telecommunication Union) grid lasers that are tuned to a specific wavelength. CWDM typically uses bands S, C and L normally with 20nm channel spacing and center wavelengths at 1491nm, 1511nm, 1531nm, 1551nm, 1571nm, 1591nm, 1611nm. CWDM wavelengths can start as low as 1270nm. The same techniques are used for DWDM networks as well. However, the wavelength spacing is much tighter at .4nm (50Ghz), .8nm (100Ghz) and 1.6nm (200Ghz) from 1525 to 1615nm which drives up component cost for both lasers and passives.

Most optical transceivers are now designed to comply with the SFP industry-standard MSA (Multi Source Agreement). Ethernet switches, routers and SONET multiplexers have all now incorporated SFP standards. Low-cost SFP repeaters can provide 3R regeneration (re-amplify, reshaping and retiming) to overcome loss and dispersion limitations of long links.

Figure 2. SFP modules dramatically improve price-performance and flexibility of optical networking.

SFP transceivers feed optical combiners and splitters called optical mux-demux passives. These low-cost devices separate wavelengths or lambdas while avoiding high insertion losses. These devices use GRIN (GRadient INdex) lenses or thin filters that are normally packaged together to filter two or more wavelengths. GRIN lenses focus light through a precisely controlled radial variation of the lens material's index of refraction from the optical axis to the edge of the lens. By gradually varying the index of refraction within the lens material, light rays can be smoothly and continually redirected towards a point of focus. This allows a GRIN lens with flat or angle polished surfaces to collimate light emitted from an optical fiber or to focus an incident beam into an optical fiber. End faces can be manufactured with an anti-reflection coating to avoid unwanted back reflections

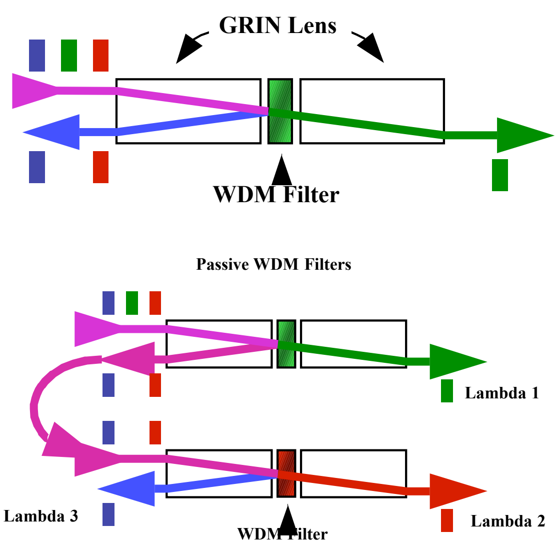

Figure 3. GRIN assures stable and reliable optical networking performance through precise lens characteristics.

Most metro networks have larger fiber counts and need fewer channels to transport a full broadcast video lineup to remote hubs. A four lambda system carrying four Gbps of bandwidth over 100 km costs under $2,500 per GigE. Two GigE links can carry over 480 standard definition video programs. Low-cost regeneration at drop sites or mid span repeaters can be deployed to overcome long distances.

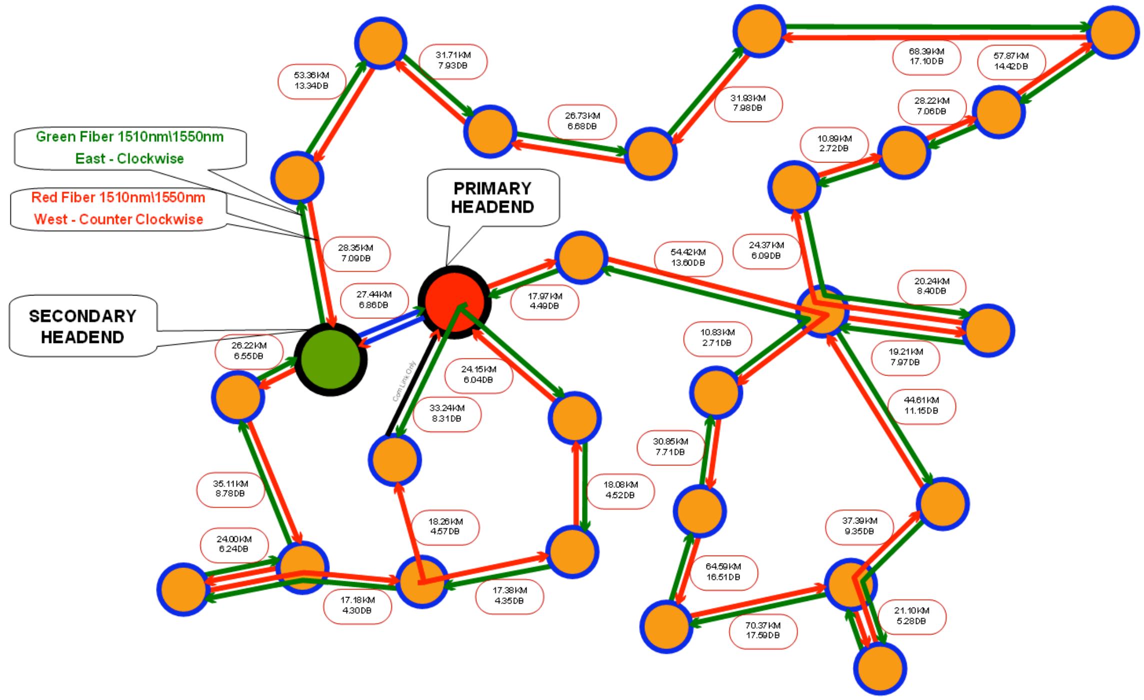

Figure 4. A passive CWDM network utilizing primary and secondary headend feeds and a bidirectional transport architecture for service assurance.

The real-world, large metro-based CWDM network represented in figure 4 provides digital video transport between two redundant headends and 24 digital hubs serving over 1.3 million homes. Two GigE links on a single fiber provide capacity for up 48 38.8Mbps MPTS broadcast video multiplexes. Each hub has dual route diverse links fed from two separate headends which prevents fiber related outages.

The biggest advantage of passive CWDM is simply the lower cost. Wide channel spacing enables less stringent laser performance requirements and lower cost optical passives. Small, integrated SFP optics require less rack space and power consumption. Passive DWDM systems are best when more that eight Gbps are required and fiber resources are limited. Passive WDM networks are very simple to provision and expand when required. In short, there are less active or moving parts to break which results in better availability. WDM systems are able to transmit multiple bit-rates and protocols, which allows the use of optimized application devices. SONET for CBR (Constant Bit Rate) T1 and DS3 traffic, Ethernet switches and routers for IP traffic and MPEG-2 routers for GigE over IP traffic allows systems to deliver more specialized features for a longer useful life.

There are a few limitations when using CWDM including the lack of wideband amplifiers. SFP CWDM repeaters work well when there are a limited number of wavelengths to retransmit but can be costly when dozens of lambdas need amplification. DWDM allows the amplification of multiple wavelengths with a single device.

ACTIVE WDM NETWORKS

Active DWDM systems work similar to passive networks but have finely tuned lasers and cross-connect features. These devices are known as transponders and muxponders.

One challenge when operating passive DWDM networks is crosstalk and balancing of optical power levels. Active DWDM networks use precise transponders for managing power levels and 1+1 optical redundant switching. High performance lasers also allow transmission over ultra-long fiber spans up to 200 km. Active transponders can also be equipped with electrical multiplexers for aggregating multiple GigE feeds into a single 10 Gbps lambda. Active DWDM networks are best when fiber resources are extremely limited and the maximum amount of bandwidth per lambda is required.

The downsides of active DWDM include added costs and more actives or moving parts to fail.

GIGE OVER SONET

SONET networks have been the workhorse of local and inter exchange carriers for over a decade. SONET was optimized for high availability or lifeline CBR transport such as DS3 and DS1 traffic. Several layers of hardware protection ensured five nines of reliability. Today's next-generation SONET network elements now have integrated lower cost layer 2 Ethernet functionality and GigE interfaces. These interfaces can be used in a one-way cross-connect mode that can efficiently transport MPEG-2 MPTS over IP to distant hubs. Figure 5 shows how one-way drop and continue cross-connects for unidirectional video traffic operate while still providing a protection path in the case of a catastrophic fiber cut.

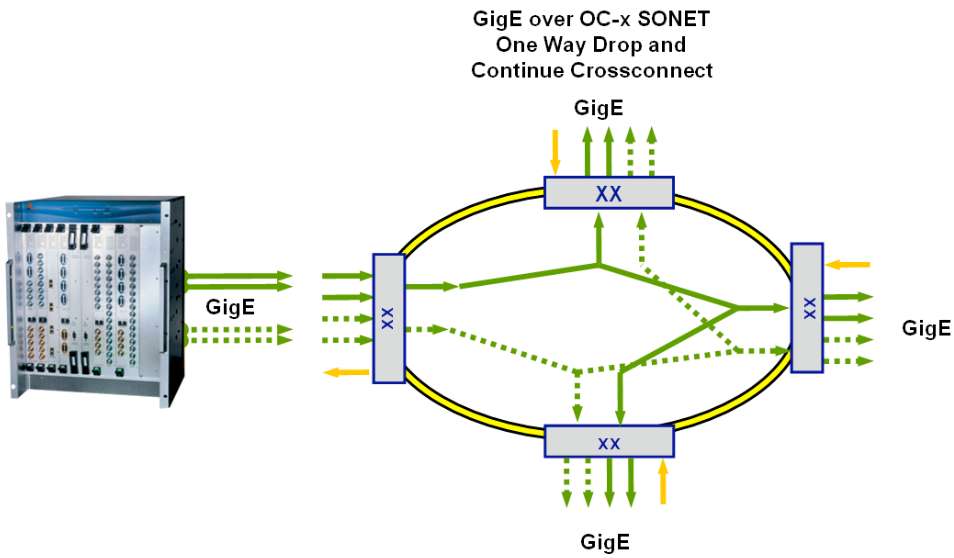

Figure 5. IP/GigE and SONET techniques can be combined for reliable and economical transport while leveraging prior investments and practices.

OC-192 SONET terminals have had significant drops in price and size over the last five years. Many operators already have IP and voice traffic traversing these rings. Excess SONET capacity can easily be provisioned for a robust MPEG-2 over IP/GigE broadcast video transport link. SONET terminals by far have the best track record for network availability performance.

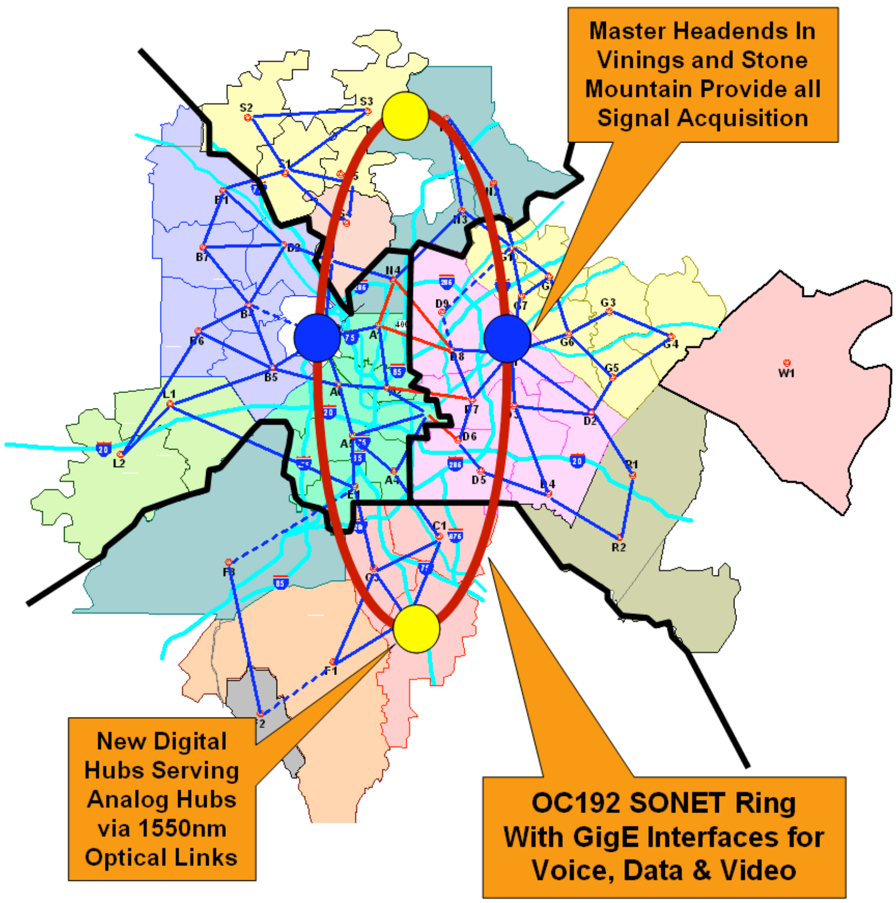

Figure 6. real-world network leveraging multi-use SONET rings.

The biggest challenge associated with SONET networks is scalability to meet the growing increases in pure IP-based traffic. Many new IP-based switching and routing platforms are catching up to the high performance track record of SONET-based solutions at far better price points. A recent SONET approach called RPR (Resilient Packet Ring) is now emerging. If prices can continue on a downward path, RPR may challenge classic IP routing with several of the advantages of packet-based routing.

CONVERGED REGIONAL AREA NETWORKS

One way to design the network is using a routed IP backbone. Cable modem and VoIP (Voice over IP) services have been running over a routed

IP network for years, while video has traditionally been sent over a parallel network. There are a number of reasons that MSOs have not been sending video traffic over an IP routed network until now:

- Until recently, IP routers did not have sufficiently robust QoS mechanisms to ensure lossless, low-jitter video delivery.
- Video equipment vendors only recently added support for IP-based video transport in their products.
- Video uses a tremendous amount of network bandwidth, and the Gigabit Ethernet capacity has been traditionally insufficient to benefit from the convergence of multiple services.

Today, however, these issues have been addressed. All major router vendors include QoS mechanisms, IP transport is rapidly becoming a standard transport option on video equipment, and 10 Gigabit Ethernet is available at very competitive price points. As a result, MSOs can now make use of the same IP routed network for all of their video, voice and data services.

IP networks can use unicast (one-to-one) or multicast (one-to-many) transmission. While VoIP and data services currently running on MSO IP backbones is unicast, video services are multicast in nature. As a result, video services carried over IP networks are typically multicast. It is critical to have a high-performance multicast-enabled IP network in order to carry video services.

A key benefit of running video services over IP routed networks is the fact that routing protocols (e.g., OSPF and BGP-4 for unicast; PIM-SM and MBGP for multicast) make the network highly dynamic and robust. Rather than having operators statically define paths from a video feed in a headend to a hub site, the IP network will automatically calculate and transmit along the least-cost routed path. If there is an equipment failure, the routers will dynamically discover the failure and re-route around it – without manual intervention.

Today's multicast IP networks generally use ASM (Any Source Multicast). In ASM networks, receivers (e.g., a splicer/groomer in a hub site) use a protocol called IGMP (Internet Group Management Protocol) version 2 to join a multicast group. It is called "any source" multicast because the IGMPv2 protocol does not specify the sender, only the group. After sending the IGMP request, the network re-configures itself to make that video feed available to the requesting device.

When using ASM for multicast IP, a multicast routing protocol is generally required to ensure that the multicast traffic is routed to all the correct network destinations; PIM-SM (defined in RFC 2362) is typically used. In PIM-SM, all traffic initially is transmitted via a router designated as the RP (Rendezvous Point). Because all multicast traffic is sent via the RP, an appropriate multicast stream can be found for each multicast group. The disadvantage of this approach is that the RP can be a bottleneck and potentially a single point of failure.

A new approach called SSM (Source Specific Multicast, RFC 3569) eliminates these challenges. With SSM, receivers must specify both the multicast group *and* the IP address of the multicast source. IGMPv2 does not support this capability, so IGMPv3 – a new version of the IGMP protocol – is required in order to support SSM functionality. A key advantage of SSM is that because the receiver specifies a source in advance, a rendezvous point is no longer needed. This eliminates a potential bottleneck and point of failure from the network. The short term challenge with SSM is that most existing devices do not yet support IGMPv3.

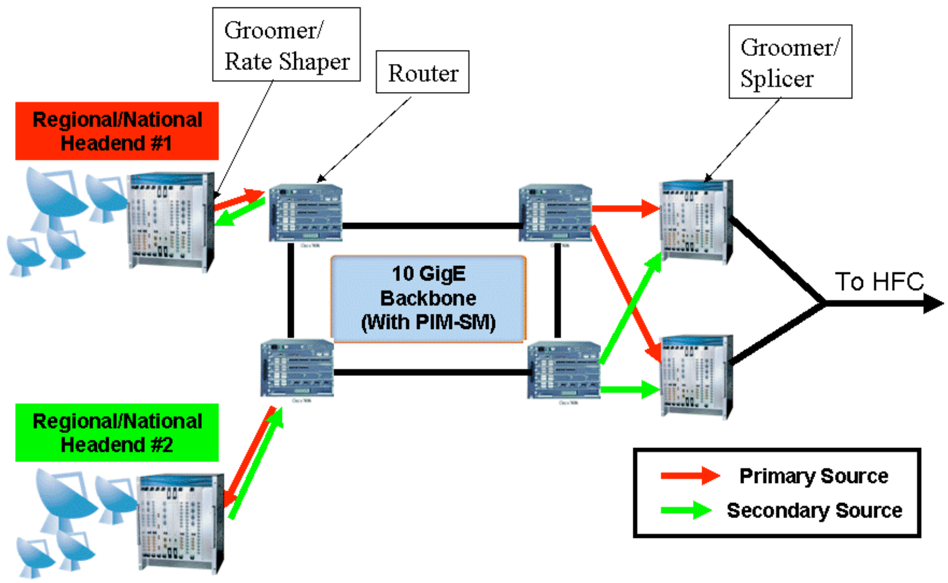

Figure 7. Incorporation of IP routers is a substantial modification of cable architectures, allowing high flexibility to support various functionalities and future directions.

An advantage of moving to an IP routed network is that the MSO can converge video and other services such as data and VoIP services on the same IP backbone. With only one network to manage, operations costs are reduced – and bandwidth can be shared among the various applications.

Running a converged network requires QoS mechanisms to be in place. Video traffic is very sensitive to jitter and packet loss, while VoIP is sensitive to latency and packet loss. In IP networks, the DiffServ protocol (defined in RFC 2474) is the most commonly used QoS mechanism. A DSCP (DiffServ CodePoint) is a six-bit field in the IP header of each packet – specifying the queueing behavior for that packet. Video and VoIP packets will be marked as high priority, and will be transmitted ahead of data packets. Using DiffServ, there can be lossless, low-latency, low-jitter transmission of video and VoIP even in the face of data congestion. Another protocol called MPLS (Multi Protocol Label Switching) can be used in conjunction with DiffServ, providing additional control over which traffic flows over which links. The disadvantage of MPLS is that it can add significant complexity to network operations.

Detractors debate the merit of converging video and other services on a single IP backbone, citing a number of concerns:

- On a converged network, video services can be brought down by denial of service attacks.
- Video programming is always-on, and of fixed bandwidth, reducing the statistical gain made possible by the sharing the network with other services.
- A converged network increases the complexity of network configuration, requiring careful QoS and traffic engineering to ensure that jitter and packet loss do not negatively impact video quality. This may increase the operating expenses significantly enough to negate any benefits of a converged network.

Several mitigating practices can address these challenges.

A primary network design consideration is security. This should include location of encryption which can be central, at the acquisition site, or at the edge, in the RAN (Regional Area Network). From a practical standpoint, each RAN must have its own DNCS (Digital Network Control System), or something similar, for handling communications with STBs in the region, and they already have CA (Conditional Access) systems in place. As such, the content need not be encrypted centrally.

However, particularly in a converged network, content owners – fearful of internet hackers somehow getting free, unencrypted programming – may require that all content be encrypted. As such, IPSec or a similar protocol may be used to encrypt traffic between the central acquisition site and regional or metropolitan headend. This remains a challenge, since IPSec devices generally do not have sufficient performance to encrypt such a large quantity of data.

Because a large number of headends (potentially every system owned by a given MSO) rely on feeds from the central acquisition headend, guaranteed resiliency is a must. To maintain such resiliency, several actions should be implemented. Programming identical to that at the primary headend must be available from a secondary headend that is in a geographically separated location. In a converged network, video traffic must be given bandwidth guarantees to ensure the network is fully non-blocking for video traffic. Video quality monitoring capabilities must be in place at each regional or metro headend. Program-level redundancy should be used, and quality monitored continuously on a per-program basis.

Ensuring video quality across a national network is not easy. In general, today's cable networks may have an end-to-end latency of 50 ms. As MSOs move to national networks, latency will increase, additional jitter will be introduced, there will be multiple paths from source to destination. In addition, there are more points at which congestion and packet loss may occur.

For MPEG traffic, the maximum acceptable jitter is 500 ns. Encapsulating

MPEG in IP alone introduces significant jitter, since up to seven MPEG packets are encapsulated in one 1,500 byte IP (over Ethernet) packet, producing queueing delays that introduce far more jitter than the specification allows. As a result, it is necessary to de-jitter the traffic at the receiving side. Poor de-jittering may result in stutter in the picture, and lower video quality.

To ensure high quality, the video portion of the network must also be virtually free of packet loss.

CONCLUSION

Passive WDM, active WDM, GigE-SONET and converged RAN are all worthwhile considerations for cable operators who want to leverage IP technologies for transport of core programming services. This has further advantages in to varying degrees enabling convergence with IP techniques used for emerging services like VOD as well as voice and data offerings. As a result, management capabilities and economics are improved. Benefits can be extended to national fiber distribution that revolutionalize program sourcing practices.

However, all of these techniques do bring their relative advantages and challenges. Selecting for short-term optimization can position an operator less well as needs change. Key to this is building current IP infrastructures on robust, flexible and scalable platforms that are programmable for ongoing implementation of the best and most recent transport techniques.

Mark Davis
Vice President,
Network Solutions
mark.davis
@bigbandnet.com

S. V. Vasudevan
Chief Architect and
VP Systems Engineering
vasu
@bigbandnet.com

David Brown
Director,
Product Marketing
david.brown
@bigbandnet.com

BigBand Networks, Inc.
475 Broadway
Redwood City, CA 94063
650 995 5000
http://www.bigbandnet.com