BROADCASTERS CAN'T MISS WITH HITLESS TECHNOLOGY

Hitless protection switching uses two separate transmission paths to deliver identical packet streams from a source to a destination. Upon arrival, any errors in one path can be completely removed through instantaneous, transparent switching from one path to the other. Due to increasing availability of equipment that supports hitless switching, coupled with declining network costs, the technology is going mainstream allowing broadcasters to deliver ultra-reliable signals from remote venues to studios and within production and distribution networks.

Introduction

While being hitless isn't a path to success for baseball batters (or aspiring rock stars), hitless signals are another thing entirely. Smooth, uninterrupted video and audio streams provide the best possible viewing and listening experiences. Signal dropouts, glitches, and other faults can be extremely distracting to audiences, so broadcasters have long been interested in technologies that can reduce or eliminate the number and intensity of signal disruptions. Preventing glitches can be hard, particularly on wide area networks (WANs). There, random packet loss and multi-packet burst losses can be encountered, in part due to resource contention and partly due to the many types of signals that flow between millions of sources and destinations every second. Hitless protection switching provides a field-proven way to reliably deliver signals across many different network topologies.

Network Errors and Broadcasting

Live television broadcasts are very demanding environments for signal transmission. Video signals, even compressed ones, use high-speed data flows that need to be delivered in an unbroken stream. Many of the communications paths used by broadcasters are one-way (such as satellite links and over-the air broadcasts). Network delays can also cause problems for some broadcast applications, such as live interviews conducted between people in two different physical locations.

All networks can occasionally experience bit errors or lost packets. Signals sent over the airwaves or through space can experience interference from many sources, including lighting bursts and cosmic rays. Signals sent locally over copper or via long-haul fiber-optic links typically

Hitless protection switching has been used to deliver streams for major sports events around the globe.

exhibit much lower error rates, but can experience occasional data losses. Data packets flowing through shared resources on public or private networks can also experience occasional congestion that causes packets to be dropped. Overall, connections using public networks can normally expect to see packet loss rates on the order of one in a thousand (0.1%). Private networks perform significantly better, but will still be impacted by occasional random or burst packet losses. While these rates may seem low, video signals used in contribution feeds often occupy thousands of packets per second. Even on private networks these signals should be expected to experience lost packets every few minutes.

Television broadcasters have long used technologies that can overcome minor signal degradations. One popular choice is forward error correction (FEC) that is widely deployed on satellite links. FEC works by adding extra data to the transmitted signal that can be used by the receiver to correct garbled bits and bytes. The number of errors that can be corrected in each block of data is governed by the amount of FEC information that is added – the greater the amount of FEC data present, the more errors that can be corrected. For example, the commonly used Reed-Solomon FEC algorithm can properly correct one bad byte for every two bytes of FEC data added to a block of bytes. In cases when the number of errors exceeds the correction capacity of the FEC system then the entire data block is lost.

Another popular error correction system called ARQ (for Automatic Repeat reQuest) uses a communication path from the receiver back to the transmitter to signal when data blocks are missing or have been received with errors. Once the transmitter receives a lost packet indication, it will re-send the bad/missing packets. This system works well for software files and other data that needs to be bit-for-bit perfect. However, the delays associated with detecting errors, sending the re-transmission request, and re-transmitting the bad packets can slow down or even disrupt high-speed video signals.

While both FEC and ARQ technologies can protect against occasional missing or corrupted data, neither solution is capable of withstanding longer term outages with dozens or hundreds of consecutive lost packets. In the case of FEC, there simply isn't enough good data to calculate the missing values. With ARQ, long outages can overflow buffers and prevent packets from being re-transmitted. For live broadcasts in particular, another solution was needed.

Hitless Protection Switching

Hitless protection switching was developed specifically to address both short-term and long-term outages that occur on wide area networks. In its most basic form, a hitless sender transmits two identical packet streams over two separate network paths. At the receiver, the two streams are re-aligned using adaptive buffers, and a single output is created using the good packets received from either one path or the other. As long as at least one copy of each packet is received from one of the two signal paths coming into the receiver, an error-free output can be produced.

This technology has a long history of successful applications, including a number of high-profile international sporting events fed to billions of live viewers around the globe. One major broadcaster said "Hitless technology allowed us to deliver streams from South America to North America without losing a single packet during live video transmission for the full duration of our multi-day event coverage." Another senior broadcast executive has indicated that hitless switching is a key enabling technology for reliable local and long-haul signal transmission as the company deploys hybrid IP/video networks today and migrates to all-IP based studio systems in the future.

SMPTE ST 2022-7 "Seamless Protection Switching of SMPTE ST 2022 IP Datagrams" provides an industry-accepted standard for implementing hitless protection switching. Publication of this standard was an important step forward, as it allowed hitless devices produced by different manufacturers to fully interoperate with each other. Interoperability testing of hitless technology was successfully performed between several manufacturers during the annual VidTrans conference hosted by the Video Services Forum industry consortium, and has also been featured at several Alliance for IP Media Systems (AIMS Alliance) interoperability demonstrations at recent broadcast industry trade shows. The packet formats defined in the newly-developed SMPTE ST 2110 standards for professional media over IP networks have also been carefully crafted to ensure compatibility with SMPTE ST 2022-7.

As shown in Figure 1, a hitless protection sender creates two (or more) packet streams with identical Real Time Protocol (RTP – defined in IETF RFC 3550) timestamps, sequence numbers, and payloads. Different UDP, IP, and Ethernet packet headers can be used on the two streams to enable them to be routed over different network segments. When the streams reach the receiver, the incoming packets are first fed into buffers to allow them to be time-aligned. This is accomplished by adjusting the data-reading pointers in the two buffers to both point to packets with the same RTP timestamps and sequence numbers. Each matching set of data packets is then retrieved simultaneously from the two buffers and then analyzed to determine if there are any errors present. Whenever missing packets are encountered, those positions in the buffer are skipped. A single output stream is created by combining the two streams and selecting the best packet available at each point in the sequence.

To obtain the maximum benefit of hitless protection switching, the two streams should be routed over completely separate networks. In other words, there should not be any common network devices or links anywhere along the path from the sender to the receiver. This restriction is important to ensure that no single point of failure can cause an interruption of both streams at

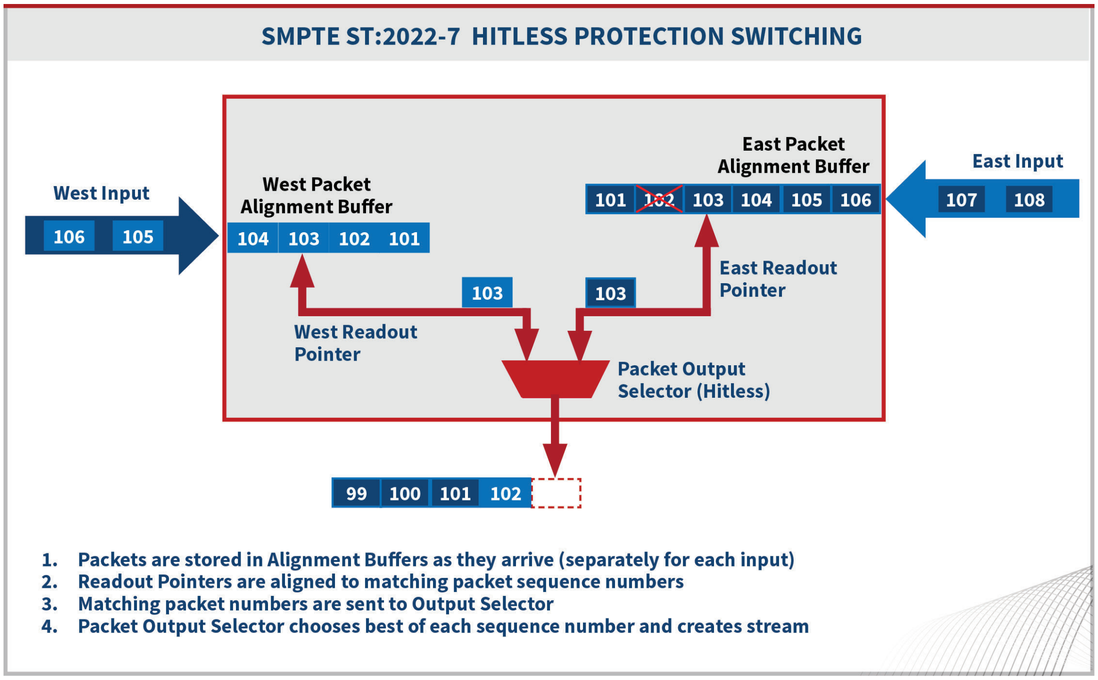

Figure 1.

the same time, which would lead to a loss of video at the system output. In practice, this can be achieved using redundant equipment within a network facility (such as a main/backup IP core) or by carefully working with WAN service providers to ensure that the two paths are not routed through the same physical location or network segment.

In any hitless application, there is a limit to the size of the buffers that are used to accommodate the difference in arrival times between the two streams. This difference, called skew in the standard, has a direct effect on the size of the packet buffers in the receiver. Receivers that can support larger amounts of skew require larger amounts of expensive high-performance memory, particularly for high-speed, uncompressed signals such as 3 Gbps 1080P video signals. Plus, if large buffers are used when they are not needed, the overall signal end-to-end delay of the signal can be increased unnecessarily. In SMPTE ST 2022-7 three receiver classes were defined and a fourth is planned in an upcoming revision. The table below summarizes the various receiver classes that have been defined.

Application areas

The traditional application for hitless protection switching is to support long haul contribution links. This was the area where the technique was first applied, and formed a lot of the background when the specifications were originally developed. Long-haul links typically include several network segments and can cross multiple carrier boundaries, so protecting signals against both short and long duration outages is a valuable capability.

As RTP-based protocols take over interconnection duties within IP studios, hitless protection switching can make sense within the walls of a production studio. Ethernet links are designed to meet a worst-case BER of 10e-12, which means that an error can occur once every 100 seconds on a 10 Gbps link and still remain within spec. A single bit error on an Ethernet link will cause an entire packet of IP traffic to be dropped; this in turn would mean that the data for several hundred pixels would be lost. Hitless protection switching would easily be able to handle these errors, and might not require adding equipment in a production facility that was constructed to provide 1:1 redundancy.

More broadcasters are using terrestrial fiber-optic connections in place of satellite links to distribute their content to major partners, such as multichannel distributors (MVPDs) and OTT program delivery systems. This trend has been driven by consolidation of MVPD head-ends into fewer, national centers; by the push for higher signal quality through fewer compression steps; and by the greater reliability offered by fiber-based systems. OTT service providers are critically dependent on highly reliable content feeds into origin servers – any content that does not make it into an origin server will not be available to OTT viewers. Hitless protection switching can be used to improve system reliability and signal availability for the critical first link to both MVPD and OTT providers. At these levels, even momentary outages can affect millions of viewers simultaneously, so ensuring system dependability is a good practice.

RECEIVER CLASS SKEWS

Class	Name	Max Skew for Low Bit Rate (less than 270 Mbps)	Max Skew for High Bit Rate (270 Mbps or higher)	Application
A	Low Skew	10 msec	10 msec	Campus/Metro
B	Moderate Skew	50 msec	50 msec	Regional
C	High Skew	150 msec	450 msec	Global

Hitless Benefits

As compared to other error correction technologies, hitless protection switching offers several benefits, including:

- Ability to work with high-speed signals, including uncompressed HD/UHD video and 10 Gbps data flows. The simple packet-by-packet decision system avoids the computational complexity of many FEC schemes, which on depend advanced mathematical formulas that are difficult to implement for high data rates.
- Can compensate for long-term outages, such as the complete loss of one leg of the network for an extended duration. The only requirement for reliable hitless operation is for both legs to not experience failure at the same time. This is why it is important to ensure that there is no common point of failure that can affect both paths between the sender and the receiver.
- Adds insignificant amount of delay, because the packet comparison mechanism can operate on the two streams simultaneously. The only buffering needed is for the packets that arrive first on one path be buffered long enough for the packets on the other path to arrive. This buffer is therefore no worse than the worst-case delay between the sender and the receiver on the two paths, which would be needed for any type of primary/backup architecture.
- Hitless protection switching works extremely well in the fully redundant configuration that is commonly used by broadcasters. In fact, it could be considered an improvement because the "Backup" link is always on and is thus continuously monitored. Hitless protection could be considered a way to enhance the long-term reliability of a fully redundant network because it eliminates the all-too-frequent problem of an unexercised backup not being ready to take over when it is most needed.
- Able to support any RTP packet stream, whether it contains video, audio, or large data transfers. As long as the two streams are sent out with identical RTP timestamps and packet sequence numbers on each matching set of packets, a hitless receiver should be able to reconstruct a single output stream, provided the skew is within limits.

Conclusion

With ever-increasing demands for transporting high-bandwidth signals at improved quality levels, hitless protection switching has become a valuable tool for broadcasters. Multiple equipment and service providers have demonstrated interoperability between devices that support SMPTE ST 2022-7. As the only available technology that can withstand partial network outages lasting multiple seconds, coupled with the flexibility to support uncompressed UHD video and 10 Gbps data flows, hitless protection switching has a "can't-miss" future in high reliability networking.

Artel's SMART Media Delivery Platform™ supports SMPTE ST 2022-7 Seamless Protection Switching for high reliability networking.

Using Hitless Links for Data

Data files of all types are a challenge for long-distance transport networks. In virtually every application, these files must be transported completely intact, without a single lost packet. The impact of single byte error in computer code can be catastrophic, and compressed media files can also suffer multiple frames of artifacts from a single bit error.

To prevent problems with file transfers, many systems use ARQ or some other form of protocol to re-send data packets that do not arrive properly at the receiver. ARQ capability is built into TCP, but is also used in a variety of proprietary file transfer acceleration protocols based on UDP. Two main factors affect ARQ performance: the circuit round trip delay, and the packet loss rate. The round trip delay is important because it affects the amount of time it takes for lost/corrupted packet messages to travel from the receiver back to the transmitter and for the packets to be re-sent. Packet loss affects how many of these messages need to be sent.

Hitless protection technology can improve packet loss rates by several orders of magnitude, and can also protect against long-duration losses on one leg of the transmission. This is a significant benefit to high-speed file transfers, as it reduces the need for buffers on either end of the link, and also dramatically reduces (and potentially eliminates) the need to re-send lost packets, thereby significantly improving file transfer speeds.

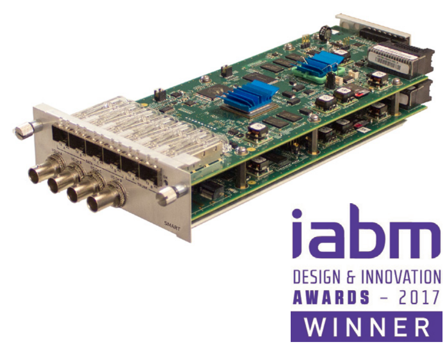

Related Products

The SMART Media Delivery Platform™ is a four-channel auto-sense SD-SDI/HD-SDI/3G over IP multi-function gateway with an integrated non-blocking Layer 2/3 switch. Winner of the IABM Design & Innovation Award 2017 in the Content and Communication Infrastructure category, the SMART Platform supports SMPTE ST 2022-7 and has been designed to attach seamlessly to the IP network without the need for external network elements. Functionality of the platform may be added or upgraded via software download. The flexibility in function capabilities along with greater port density drastically reduces power consumption, reduces size and cost, and allows for a seamless transition as the end user's network evolves. is an employee-owned business. Learn more about the SMART Media Delivery Platform at www.artel.com/SMART

About Artel

Artel Video Systems is a world-class provider of innovative, real-time, multimedia delivery solutions serving global markets. Today the majority of the live events in the US traverse Artel products to support their mission critical work-flows. Artel's expertise in IP- and fiber-based technologies spans more than 30 years and has established Artel as a trusted partner in the development of reliable, standards-based, IP infrastructures. Artel's integrated solutions include fiber and IP based multimedia delivery, data networking, IP streaming, and precision timing. Artel is an employee-owned business. More information is available at www.artel.com.

World Headquarters
5B Lyberty Way
Westford, MA 01886
USA
Tel: (978) 263-5775
Fax: (978) 263-9755
Artel.com | info@artel.com

Cover photo by Arturo Pardavila III (Lorenzo Cain swings at a pitch) [CC BY 2.0 via Wikimedia Commons]

September 2017

AR207-900001-00_B