Core Insights

From the experts at The Broadcast Bridge.

TDM Mesh Networks - A Simple Alternative To Leaf-Spine ST2110

THE BROADCAST BRIDGE
original independant content

supported by: RIEDEL

Introduction by Tony Orme - Editor at The Broadcast Bridge.

IP is delivering unprecedented flexibility and scalability for broadcasters. But there is a price to pay for these benefits, namely, the complexity of the system increases significantly as we add more video and audio over IP.

The principal reason for this increased complexity lies in timing and synchronization. We may have moved on from line and field syncs, but their presence still influences broadcast workflows and video and audio processing.

Although timing reference signals may be hidden in SDI and their influence isn't particularly evident, everything changes when we move to ST2110.

PTP has its history in industry, has been used for many years and is well understood in the IT community. However, it's still relatively new for broadcasters and we're still ironing out all the wrinkles, not just in our understanding, but also in how we apply it to the specialist media streams we work with.

Video and audio streams are quite unique as the data they create is relentlessly constant, especially when working in the uncompressed domain. This leads to tight time tolerances making interfacing the synchronous SDI and AES transport streams to IP packet networks even more of a challenge.

Building, installing, monitoring, and understanding IP networks for video and audio transport is not for the faint hearted. It most definitely delivers impressive benefits, but there are occasions when these benefits are just not needed. One example is that of this years European Song Contest (ESC).

The limited working environments, rigging and test availability, and rehearsal times are not the exclusive to the ESC, but are instead typical of ad-hoc international events the world over. IP would lend itself well to such a use-case, but a purely IP solution proves difficult to manage and implement.

Another solution is available through the use of TDM (Time Domain Multiplexing). Based on a mesh network topology, TDM provides many of the benefits of IP without having to get bogged down in the detail of operating it. Switching nodes are connected through fiber optic cabling to provide high-speed interconnectivity, flexibility, scalability, and resilience.

Each of the nodes not only acts as a router between network connections, but provides interface gateways to all the broadcast formats such as SDI, AES and analog audio. Furthermore, IP interfaces are available to provide IP routes across the network without having to run a separate IP network.

A software management and monitoring system makes configuration and routing of devices easy as well as providing resilient connectivity that self-heals should a fault occur.

TDM has another benefit that may not be immediately obvious, that is, very low and predictable latency. The synchronous nature of time-data slots within the TDMs transport stream frame means there is always availability for the data being transported. There's no congestion and if a link runs the risk of becoming oversubscribed, extra fiber cabling and nodes can be easily added to scale the infrastructure accordingly.

Bringing together SDI, AES, analog and IP all into a fully distributed mesh network is another option available to broadcasters. Complete IP infrastructures certainly have their place, but TDM not only offers a viable alternative when the complexity is not required, but also provides the ultimate in hybrid systems.

Part 1 - Balancing Technical Requirements

IP is well known and appreciated for its flexibility, scalability, and resilience. But there are times when the learning curve and installation challenges a complete ST-2110 infrastructure provides are just too great.

Many broadcasters are moving to IP because of the immense flexibility it offers, especially when we consider the availability of 4K, HDR, WCG and higher frame rates. Distribution of multiple audio streams lends itself well to IP and many other services such as intercom, monitoring, and control can be easily added.

The principle challenge we have with distributing low latency video and audio over IP is attributed to configuring multicast streams and creating an out-of-band timing plane. From the earliest days of broadcasting the synchronizing information was embedded in the video signal and this system continued through to SDI. It's only when ST-2110 was introduced did we separate the video and audio essence from in-band timing.

Scalable Resource

Flexibility and scalability lend themselves well to IP infrastructures. The IT industry demonstrates this on a daily basis through network scaling and resource expansion. When COTS methodologies are employed it's relatively easy to increase server resource and network capacity. Software defined networks and workflows facilitate this especially if cloud systems are used.

Both video and audio streams are inherently synchronous. This isn't just about maintaining lip-sync but guaranteeing the number of video frames and audio samples created by the source devices are played back in time by the display and loudspeaker systems. If a microphone is creating audio samples faster than a loudspeaker can decode and play them then buffer overflow will occur leading to horrible audio distortion. Underflow occurs when not enough samples are created resulting in a similar distortion.

Synchronization Matters

Although the original reasons for maintaining video synchronization may have disappeared into the history books, the principle is still maintained. That is, every video frame and audio sample created, must be played back synchronously for the viewer. And this is where our challenges for timing emanate.

PTP as an out-of-band timing system provides a reference that allows video frames and audio samples to be uniformly created and synchronously viewed. The time-invariant nature of signal sampling cannot be underestimated as any variance in the sample rate will also result in non-linear distortion that is difficult to correct.

Timing systems are well understood in

industry and broadcasters have adopted PTP with its wealth of experience already well documented. However, this is a highly complex protocol with many intricacies. Monitoring isn't as straightforward as checking sync pulses on an oscilloscope, and synchronizingdisparate systems requires a deep understanding of networks and their topologies.

Solutions For Applications

That said, ST-2110 with PTP is continuing to gain a lot of interest and broadcasters are finding applications where it excels. Predominantly, these tend to be fixed installations where IT and broadcast teams come together to design, install, and maintain IP infrastructures. OB's do use ST-2110 and IP as its adoption greatly reduces the amount of cabling that results in reduced axle weight. But again, these tend to be components of much larger infrastructures where support teams have had time to understand and test this new technology.

There is more to IP than distributing video and audio; it also includes transport of control and monitoring information. For example, live studio and OB cameras need OCPs for shading, line-up, and monitoring, and these often connect to the CCU using IP. System control, routing and tally all benefit from IP as it saves on having to install custom cabling for standards such as RS422. Not only is this difficult, but custom routing systems are needed to switch control signals, further adding to the complexity and weight. SDI is a relatively static system and works with a limited number of formats, and SMPTE have done a fantastic job of keeping up with the latest broadcast standards to distribute signals. For example, 4K120P can be distributed over 12G SDI cabling along with HLG HDR and WCG. However, any non-broadcast formats are difficult to transfer.

Fiber Distribution

Limited to one service per cable, and sometimes less when dual and quad SDI are used, SDI uses a massive amount of cable that is heavy, cumbersome, and difficult to maintain. It's always the center cable tie-wrapped in a large bundle that breaks. As IP is a packet switched network the topology is completely different from circuit switched systems such as SDI, facilitating many streams to be distributed over single cables, whether fiber or twisted pair.

Cable management is a concern for many remote installations. This may not be the case for fixed broadcast facilities, but even these have limited capacity for cabling especially when considering underfloor and elevator capacity.

As in all things engineering there is always a compromise and balance to be struck. SDI may be static but is easy to use, whereas IP and ST-2110 is incredibly flexible but difficult to configure and maintain.

TDM - A Third Way

Until recently broadcasters have had to make a stark choice between SDI and IP for their installations. Balancing the requirements, advantages, and disadvantages, is often a difficult task. However, a further compromise is now available, that is TDM (Time Division Multiplexing).

TDM systems tend to be vendor specific, and this is one of their greatest strengths as managed networks simplify designs. Multiple signal types are integrated into aggregators that multiplex them onto single fiber channels. This provides all the advantages of SDI simplicity with the flexibility of IP networks and distribution, but with much of the complexity taken away.

Furthermore, cable management becomes much easier as the aggregators join through fiber connections to make a simplified network. Instead of having to run a new cable between two trucks or a truck and the arena for each signal, a new stream within the network is established through the management interface to provide virtual circuits. And the aggregators can also process video and audio so that they act as signal processors. This includes frame synchronization, HD to SD and standards conversion.

In the next two articles we will take a deeper look at TDM and the associated networks to understand how they benefit broadcasters.

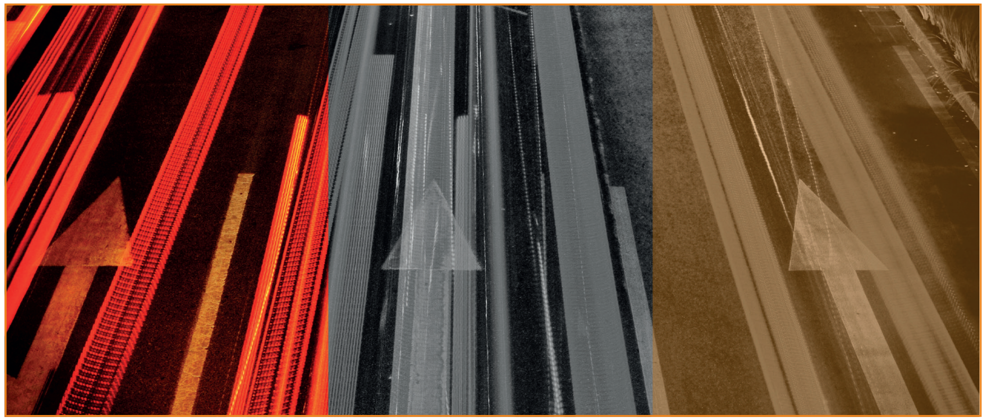

Part 2 - Knowing What To Choose

Broadcasters are no longer faced with the binary choice of going down the SDI or IP routes. The hybrid method of using TDM (Time Domain Multiplexing) combines the advantages of distributed networks with IP and SDI to deliver a fully integrated solution that helps broadcasters working across multiple technologies.

In recent years, the fiber optical backbone speeds have increased at a massive rate. It was only in 2010 that 40Gb and 100Gb ethernet speeds were considered state-of-the-art, by 2020 we now have 400Gb with 800Gb, and 1.6Tb is predicted for 2030. These speeds have been empowered by advances in fiber optic cable design and wavelength division multiplexing (WDM) technologies.

At these bitrates, the once extraordinarily high broadcast speeds have almost disappeared into the history books. With progressive HD signals using 3G-SDI, the bitrate needed is just short of 3Gbps, and 4K requiring just short of 12Gbps. This allows a staggering 130 progressive HD signals to be transferred over a 400G backbone or 30 progressive HD signals over a 100G backbone. But it gets even better as IT networks are often full duplex thus allowing for the same number of video feeds to be sent in both directions simultaneously.

IT Is Delivering

The point is IT networks now surpass the sort of network speeds needed by the broadcast industry. At least in point-to-point configurations. Although possible, it's unlikely that a studio or stadium will have 130 cameras, and this leaves open the option of using alternative topologies to the traditional broadcast centralized configurations.

Since the inception of IP broadcasters have generally been thinking in terms of spine-leaf architectures. These have the advantage of being predictable in terms of latency and have a relatively good level of resilience. However, they are hierarchical and congestion points soon manifest themselves if careful system design isn't employed.

Mesh networks differ from the spine-leaf architectures as they can be fully connected, dynamic and non-hierarchical. Each node within the mesh acts as a switch or router and is aware of the other nodes within the network. They have the ability to exchange routing information and if each node has more than one network connection the network can self-heal if a fault occurs.

Synchronous Distribution

TDM is a synchronous method of distributing data across network links. It not only has the advantage of employing fixed latencies but is predictable in terms of its capacity and signal throughput. Unlike packet switched networks, every data packet has an allocated slot in the TDM frame so buffer utilization is kept to a minimum at the receiver.

Through high-speed switching fabrics, two network connections can have their data slots in their frames exchanged, this in turn provides a routing system. For example, if a mesh node has two network connections NET1 and NET2, and a video service is using TDM slot 10 in the NET1 input feed, the node will take the data from this slot and insert it into a free slot within the NET2 connection, thus providing routing. This methodology can be carried out for any type of service as the node is switching application agnostic data within a time frame between synchronous networks.

Although the nodes are aware of each other within the overall network, it is normal practice to provide a method of management to facilitate routing and fault detection. Software defined networking lends itself well to this application and an entire mesh network can be built using it.

Reducing Congestion

The non-hierarchical nature of synchronous mesh networks means that there is less possibility of congestion. It's possible a network may reach its full capacity as all the data slots within a TDM frame could be utilized, but this is highly predictable and well within the control of the SDN management software.

Spin-leaf networks also use SDN type architectures but maintaining low latencies and reducing the risk of over oversubscription within the switches and networks is much more complex. This is mainly due to the monitoring methods adopted by the SDN system. Network switches use systems such as OpenFlow to gather monitoring statistics on their ports including datarates, packet loss and jitter. Although these are often used, their accuracy is a function of the switches ability to measure and report.

Gathering these statistics is often the role of the SDN and it uses this data to optimize the network, provide routing and predict oversubscription and latency issues. However, by virtue of the fact that the monitoring information has been integrated, passed to the SDN servers and then processed, the control is always reactionary. Very occasionally predictive systems can be used to announce the possibility of congestion etc., but there is

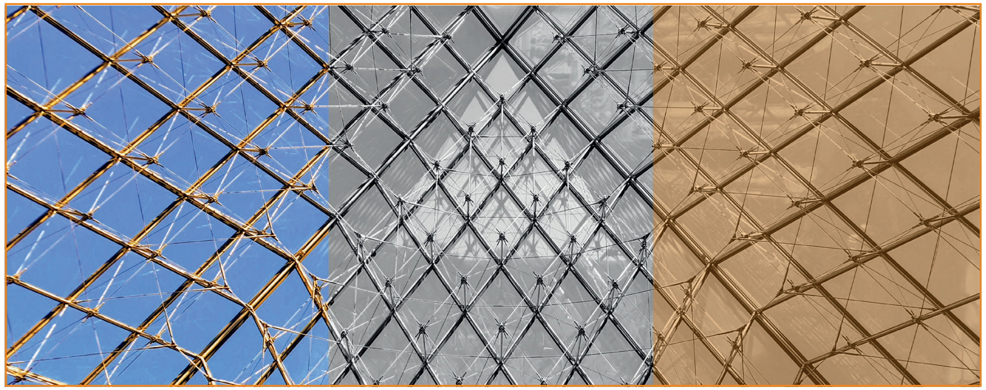

always a compromise with such systems and false positives can easily occur.

Comparing Leaf-Spine To Mesh

If we contrast the leaf-spine and mesh topologies, it's clear that the mesh networks are much more predictable in their latency and subscription capacity. The very nature of the time domain multiplexing means that every data packet from an SDI, AES or control channel, always has an allocated slot within the TDM connection between nodes. The SDN managing the network knows exactly how much capacity is available and where the data packets are within their journey.

Just as an IP packet has no knowledge of the data it is carrying and nor does it care, the same is true for TDM mesh networks. As far as the transport stream is concerned, the data slots are just transporting application agnostic data from one node to another. It's most probable that the higher level SDN software is aware of the data service, for example an HD video service, and this service autonomy is what makes TDM incredibly flexible for broadcasters.

The mesh network can easily transport video, audio, control data and even IP and ethernet packets. The non-hierarchical nature of the network topology lends itself well to self-healing systems and the synchronous time domain multiplexing keeps latencies predictably low and network congestion at an absolute minimum.

Mesh networks employing TDM work well for localized networks, but they can be routed through gateways to expand the LAN to WANs. In the next article in this series, we look at the Eurovision Song Contest as an application of TDM Mesh Networks.

Part 3 - Applications – Eurovision Song Contest

With over 4000 signals to distribute, transfer and route, the Eurovision Song Contest (ESC) proved to be this year's showpiece for Riedel's TDM based distributed mesh networked system MediorNet. Understanding the intricacies of such an event is key to realizing why TDM is such a powerful solution.

Video signals are known for using high bandwidth and bulky cables. It doesn't take too many video feeds for a bundle of cables to become bulky and heavy. Then we add the audio; although the bandwidths are much lower than video, the cables are equally as bulky. The network connections come next for the myriad of computers, OCPs, and control and monitoring units. And most importantly, the audio and control cables for the intercom, without which, nobody would be able to work.

Running these cables throughout a custom-built TV station is a difficult and time-consuming at the best of times, especially when we start running cables through the elevator conduits to move between floors. But running this number of cables through an auditorium with very little time is a daunting task. It wouldn't be so bad if the TV riggers and engineers were working in isolation, but the ESC required a multitude of crews to create the sets, rig the lights and configure the auditorium, and that's before the production teams arrived.

Pushing Technical Boundaries

ESC is not only renowned for its showcase glitz and glamour, but also for its technical excellence and for pushing the boundaries. First broadcast in 1956 to push broadcast technology to its limits, ESC helped align many European countries so they could exchange live video and audio signals, a solution unique for its time.

If instead of running hundreds of video, audio and network cables, single fibers could be used, many of the problems of rigging would be instantly fixed. Complexity, resilience, and costs would be fixed in an instant. In effect, this is what Riedel's MediorNet has achieved.

In the previous article in this series, we discussed the benefits of mesh networks and how they assist broadcasters in localized network settings. Ad hoc auditorium events are generally self-contained with either a production crew on site or the camera and audio sources being streamed back to a centralized control room.

Signal Requirements

This year's ESC production in the Ahoy Arena, Netherlands, needed 24 cameras, 140 intercom panels, 75 microphones, 5 mobile units, a technical operations center, 3 support vehicles, and 60 EVS channels. A total of 130 video feeds were required. If we just look at the cameras alone, each camera will need multiple return feeds for the camera operator, prompter feeds, tally, intercom, and program audio. Although most of these signals will be presented to the camera's CCU they still need to be routed to it.

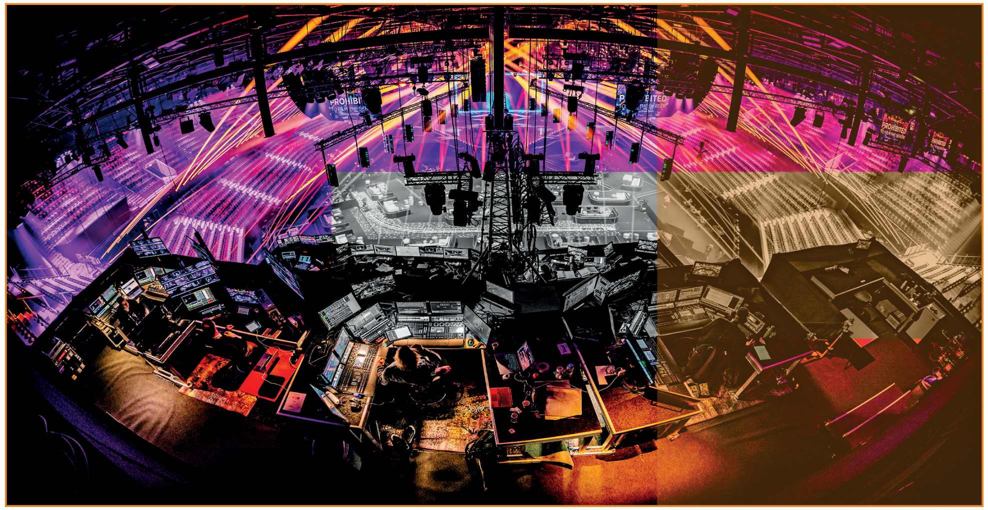

With each MediorNet device acting as both a node and signal interface, network fiber cables provided the connectivity between them. This facilitated the connections for SDI video, AES audio, control, tally, intercom and computer connectivity into the network. The management software not only provided an instantaneous health check of the entire network, but also allowed ad-hoc configuration after the initial installation.

Changing Production Requirements

Anybody who has worked on any outside broadcast, no matter how well planned, will know that things always change after the installation has been completed. Extra monitoring positions, intercom panels, and computer points are always needed. And if we factor into this the forty countries contributing and receiving audio and video for the event, then a completely flexible system is needed.

The major advantage of the distributed network is that extra feeds, receivers and monitoring streams can be added anywhere within the network at any time. Admittedly, enough capacity will need to be provisioned into the infrastructure, but even if it's not then a single fiber and node can be installed with minimal interruption. Adding an extra video service requires connecting the physical device to the node and then configuring it on the software manager.

Despite the resurgence in streaming video and audio over IP, much of the equipment currently available in broadcasting still uses traditional SDI and AES interfacing. Although interface units to facilitate connection to IP exist, configuring full ST-2110 systems for ad-hoc events is time consuming and complex.

Aggregating Signals

The Riedel MediorNet system fixes this by providing SDI and AES connectivity into the TDM distribution system. Large capacity fibers with up to 400Gbps connectivity, easily facilitate the SDI and AES integration. Once the feed is in the network, the software interface is used to route the signals anywhere within the network. The simplicity of this cannot be overstated as the network autoconfigures to make the contributing signals available anywhere on the network.

Commentators often require last minute feeds and running cables to them at an event such as the ESC is impractical and potentially dangerous. With adequate cabling infrastructure, and with a full knowledge of all the network, the software management system was used to route new video, audio and even comms circuits to the commentators.

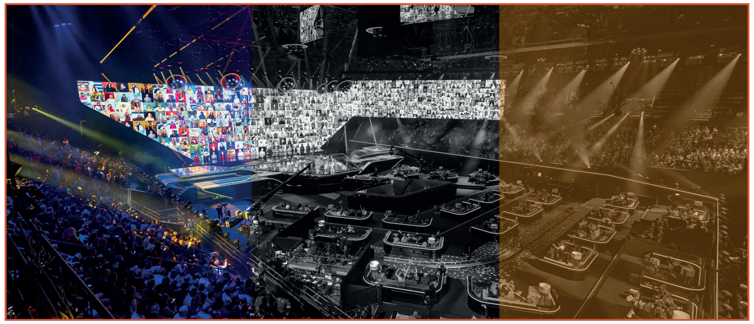

The commentary subsystem of the TDM infrastructure was complemented by an IP-based network built on the MediorNet IP product range. Inside the commentary control room, the FusioN standalone IP processing devices provided 16-PiP ST-2110 multiviewers, while the commentary booths relied on Riedel's intelligent MuoN SFPs for bi-directional gateways from MADI to ST-2110.

Monitoring is always a challenge for ad-hoc events, and this was no different at ESC. With such a complex network containing in excess of fifty nodes and interface units, knowing where a system is experiencing problems is critical. Cabling in outside broadcast type events is always more difficult due to the temporary rigging that increases the susceptibility to damage. Being able to detect cable faults and automatically re-route signals is critical for today's broadcasters, especially as we consider the ongoing effects of Covid-19.

Reliable Low Latency Delivery

As the TDM is synchronous and data packets are multiplexed into well-defined timeslots, latency becomes less of an issue as there is no contest on the network and therefore buffer sizes can be kept to a minimum. The distance between nodes is well known and the use of fiber means the transport delay is insignificant. Consequently, the whole system delivers low latency with incredibly high levels of reliability.

Delivering to over 183 million viewers across 36 markets demands the ultimate in resilience and low latency. Advanced monitoring methods and fault detection meant the MediorNet was the perfect solution for an ad-hoc event such as the Eurovision Song Contest.

This Core Insights document was written by the Editorial team at The Broadcast Bridge in collaboration with our sponsor Riedel Communications.

The Broadcast Bridge publishes hundreds of free articles like this every year. Subscribe online and create your own personalised home page and email feed.

www.thebroadcastbridge.com