E2X Bandwidth and Bit Error Requirements for Ethernet Synchronization

Introducing a Reliable Real-Time Point-to-Multipoint E2X Transport Protocol

PHILIPP SCHMID
Nautel Limited
Bangor, Maine

INTRODUCTION

The introduction of the generation 3 digital radio systems broadcast architecture transforms the broadcast system from an audio based broadcast system to a general digital data broadcast system that carries audio content along with advanced application services (AAS). The transition to this architecture necessitates changing traditional audio based studio transmitter links (STLs) to generic IP based data links.

Whereas IP streams are ideal for transferring arbitrary digital information not limited to audio content, many of the challenges inherent to IP based streaming are now imposed on this architecture. The nature of this system places stringent requirements on the STL that can cause significant on air data outages.

The network and bandwidth requirements of carrying the data stream to the exciter has already been well analyzed and published[1]. However, to deploy a system that is successfully synchronized based on the data stream rather than external GPS based synchronization brings an additional set of challenges which are presented in this article.

An alternate transport protocol to UDP or TCP/IP is presented in this article, that drastically improves data loss across the STL by implementing an automatic repeat request (ARQ) scheme. This protocol considers the impact on studio transmitter synchronization due to retransmissions.

The protocol operation is detailed and illustrates its capability of addressing multiple exciters with a reliable data stream. The article demonstrates a main/standby configuration, a dual station topology, as well as, the protocol's applicability for satellite distribution of the E2X data stream.

GENERATION 3 IBOC BROADCAST SYSTEM

The 3rd generation IBOC radio broadcast system architecture differs from previous generations in that most of the IBOC functions and equipment are moved from the transmitter site to the studio site. Illustration 1 provides a basic overview of the architecture.

The only function in the IBOC system that remains at the transmitter site is the digital IBOC modulation performed inside the exciter engine (Exgine) modulator board. Keeping the modulation function at the transmitter site is necessary since the digital bit stream to be modulated can be transferred in fewer than 150 kbps. IBOC modulation turns this bit stream into a discrete time signal that requires the equivalent of 23 Mbps to be transferred across an STL. Though required to contain the complete FM baseband frequency spectrum, it is inefficient to carry this amount of data across bandwidth limited STLs.

The transmitter site receives two data feeds from the studio: an audio feed for traditional FM broadcast, which may be carried in any digital or analog industry standard, and an Ethernet based digital data stream for IBOC modulation. This digital data stream is not limited to carrying audio data and may contain program service data or other data services along with one or more additional compressed audio streams.

The major IBOC system components at the studio site are the Exporter, Importer and system synchronization unit. The Exporter's responsibilities include delaying the analog audio to match the digital on-air delay and passing the audio signal to the STL. On the digital side

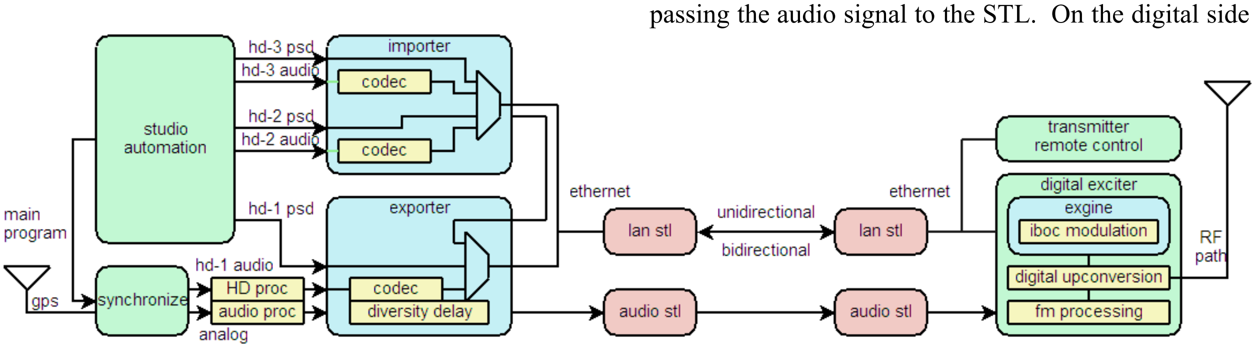

Illustration 1: Generation 3 IBOC FM Radio Broadcast Architecture

the Exporter performs the audio coding of the digital main program and multiplexing all digital content including program service data for the main program, as well as, all digital content originating from the Importer.

The Importer can capture two additional audio feeds for secondary program services which are encoded in the Importer and are passed to the Exporter along with their corresponding program service data.

The broadcast architecture is synchronous in nature in that a discrete number of audio samples captured at the studio translates into a discrete number of digital symbols produced at the exciter. The synchronization unit synchronizes the incoming audio feed to a GPS based timing reference, which in turn affects the timing of all operations in the system.

The STL link must now be able to carry Ethernet/IP based traffic along with traditional audio streams, requiring a digital STL. Many stations have already transitioned to digital STLs in favor of a digital audio feed over an analog feed for improved audio quality. Most STL vendors now provide upgrade paths to share the existing digital STL bandwidth between audio paths and Ethernet based paths. However, since most STLs have traditionally been unidirectional in nature, many STLs today also only provide a unidirectional Ethernet path. Even basic network functions, such as the address resolution protocol, often require a bidirectional data path between a sender and receiver on a network. It is important that the digital data stream protocol does not require a bidirectional data path, but it may make use of a bidirectional link if one is present.

Classifying the bandwidth requirements of the data portion of the digital data stream is relatively straight forward and iBiquity Digital Corporation has sufficiently detailed these requirements. The bandwidth requirements for the synchronization stream, on the other hand, are more difficult to determine.

SYNCHRONIZING THE IBOC DATA STREAM

The digital data stream across the STL is a generic bit stream rather than a discrete time signal. This means that this data stream cannot be sample rate converted, since each bit must remain unaltered. Consequently, both the Exporter at the studio and the Exgine at the transmitter must remain frequency locked in order to process the data at precisely the same rate. At the transmitter site the reception of the data stream is closely coupled to the digital modulation stage and, therefore, the data stream indirectly dictates the digital symbol rate. Disciplining the exciter with a second GPS receiver at the transmitter site does allow for effective frequency synchronization and relieves STL requirements somewhat. However, this approach increases the overall component count and does not control digital signal throughput delay without absolute frame alignment.

This synchronization requirement fundamentally differs from standard network communications, where data exchanges are often flow controlled by the receiver as in the case of TCP/IP. The synchronization unit cannot be placed at the transmitter site, as it does require a bidirectional STL in order to rate control the studio. Placing the synchronization unit at the studio does account for unidirectional STLs and allows for the extension of the architecture to synchronize multiple transmitter sites simultaneously.

Illustration 2 depicts a source synchronous clocking scheme operating across an asynchronous transport medium, such as an Ethernet or IP based link, and recovers the source processing rate in order to rate lock the digital modulation process. The system is disciplined by a high quality reference frequency, such as the 1 Hz signal generated by a GPS receiver, all other processes fall into place with respect to this reference signal. The audio feed is re-sampled with the GPS disciplined sample rate to ensure frequency stability rather than depending on the clock signal within the audio feed. The re-sampled audio feed is encoded and other digital processing is applied to create the data stream to traverse the link. After a fixed number of audio samples, clock packets are inserted into the data stream.

At the transmitter site, the exciter uses the clock packets to

a) recover the studio's processing rate by disciplining a frequency locked loop or a phase locked loop by aggressively rejecting high frequency clock packet jitter introduced by the link, and

b) maintain a constant throughput delay by estimating an appropriate starting time based on averaging an initial number of clock

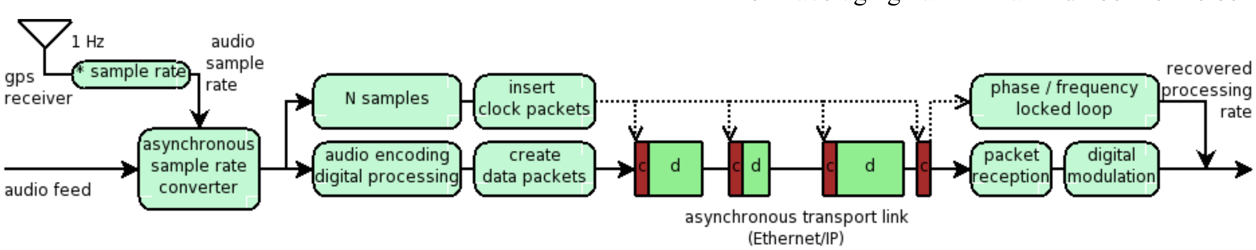

Illustration 2: Source Synchronous Clocking across an Asynchronous Data Link

packets. Clock packet jitter normally falls in the range of milliseconds, but in extraordinary circumstances can be as high as several seconds or more depending on the situation and type of STL. Therefore, it is not prudent to determine the initial digital modulation based on a single link dependent event. A phase locked loop can correct for this initial phase offset error, but a frequency locked loop will maintain this error indefinitely.

Exciter synchronization does not generally care about the link throughput delay. It will, however, react to long term changes in throughput delay, hence, STLs that automatically adjust modulation schemes dynamically based on transmit conditions are generally undesirable in this situation.

It does make sense to discipline exciter synchronization using clock packets. This minimizes the impact of various link characteristics by reducing the impact of packet serialization delays through a smaller packet structure. However, in order for clock packets to represent a true measure of throughput delay, the intermediate data packets must not impact clock packet timing, so we can determine a dedicated link bandwidth rate of

$$dedicated\ link\,rate = \frac{8 * max(data\ bytes)}{clock\ packet\ period}\,bps$$

Two synchronization levels are specified by the IBOC standard that define the maximum allowed symbol clock frequency error.

- Level I: must meet a maximum error of 0.01 ppm This requirement is difficult to attain given the typical amount of clock packet jitter even in optimal network deployments. In this case, a second GPS synchronizer unit is required at the transmitter site.
- Level II: must meet a maximum error of 1 ppm This requirement can be met using Ethernet based synchronization, provided the STL does not inject clock packet jitter greater than 15 ms and the STL bandwidth meets the minimum dedicated link rate outlined above.

The impact of data stream induced timing errors due to an insufficient bandwidth can be mitigated to some extent, however, the embedded timing information is degraded leading to less stable digital symbol rates and less predictable throughput delays.

EXPORTER TO EXCITER PROTOCOL

The data stream protocol between the Exporter and exciter components is termed the E2X protocol and is defined as part of the standard IBOC Radio System Broadcast Architecture. It implements a source synchronous clocking scheme similar to the one described in the previous section. The FM IBOC system transfers all its payload data in Layer 1 frame intervals of 1.48 seconds. This frame is broken into 16 data packets sent after a clock packet sent every 92.8 ms. However, the payload is not evenly distributed across all 16 data packets in the L1 frame as shown in Illustration 3.

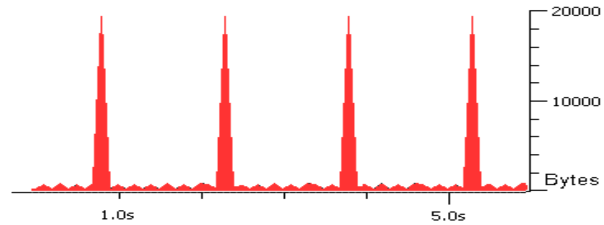

Illustration 3: Service Mode 3 Bandwidth Utilization

With a maximum data packet size ranging from 18788 Bytes to 19330 Bytes, the minimum dedicated link bandwidth is around 1.5 Mbps using the relationship developed in the preceding section. Even though these data packets are broken into smaller IP packets they are all delivered to the STL at the same time creating a temporary congestion condition. On a 256 kbps link, this congestion period can be around 600 ms affecting several subsequent clock periods. The exciter's receive buffer must be sufficiently deep to absorb this data imbalance; a receive buffer of at least one layer 1 frame of 1.48s is recommended, bringing the overall analog-to-digital delay to just under 9 seconds.

It should be noted that the E2X protocol can be successfully deployed with the recommended bandwidth requirements stated by iBiquity Digital Corporation for the E2X protocol[1], but unless a bandwidth greater than or equal to the ideal dedicated bandwidth is employed exciter synchronization is degraded using Ethernet based synchronization. GPS based synchronization can be used for STLs with lower bandwidth.

The E2X protocol follows the standard Open Systems Interconnection (OSI) networking model and limits itself to the application, presentation, and session layers. It assumes standard Ethernet technology for the physical and data link layers, but intermediate links do not strictly have to be Ethernet. The Internet protocol (IP) is employed for the network layer. E2X initially only supported the User Datagram Protocol (UDP) as a transport layer protocol, which is a fire-and-forget best effort protocol that does not guarantee data delivery. UDP can, however, address multiple destinations at once through the use of broadcast or multicast communications and also works on unidirectional STLs. The transmission control protocol (TCP) is in the process of being adopted by the E2X protocol, but at the time of writing is not yet fully supported. TCP does offer guaranteed ordered end-to-end data delivery, but is limited to point-to-point communications and requires a bidirectional link.

Since UDP is widely used in many E2X deployments, it is of interest to investigate the associated STL quality requirements. A single bit error on the STL can cause an entire data packet to be dropped, which affects an entire L1 frame in the IBOC data stream causing a 1.48 second audible HD outage. The main program may blend back to the analog audio transmission, but secondary programs and data services experience an irrecoverable interruption. Bit error rate is not an adequate measure for quantifying STL requirements in a networking environment, since many types of packet loss in a network are not directly related to bit errors. However, bit error rate is a convenient measure to assess protocol performance and provides a guideline for an appropriate signal to noise ratio on an RF based STL.

The effect of a single bit error on the STL can be contrasted between a comparable AES audio feed and the E2X data stream. With the exception of the impact on the control and status flow of the AES stream, a single bit error only extends to a single sample and even in this case errors can be concealed by interpolating neighboring samples. So it is unlikely that a single bit error is even perceivable to a listener. In the IBOC system an interruption of 1.48 seconds is very perceptible. If we define the delivered quality of service (QoS) as the ratio of flawless transmission over the total transmission time, then we can define the allowable mean time between failures (MTBF) as:

$$MTBF = \frac{error\,duration}{1 - QoS}$$

$$MTBF = \frac{1.48\,s}{1 - 0.99999} = 148000\,s$$

The required bit error rate to sustain this quality of service is then:

$$BER = \frac{1}{MTBF * bit\,rate}$$

$$BER = \frac{1}{148000\,s * 115503\,bps} = 5.85 * 10^{-11}$$

Correspondingly for AES at 44.1kHz, the same quality of service requires a bit error rate of:

$$MTBF = \frac{22.68\,us}{1 - 0.99999} = 2.268\,s$$

$$BER = \frac{1}{2.268\,s * 1.41\,Mbps} = 3.12 * 10^{-7}$$

For the digital performance this means an interruption of service roughly every 41.1 hours. The bit error rate can be converted into a packet loss rate of about $4.22*10^{-7}$. Ibiquity recommends a maximum packet loss rate of 10^{-5}, which may cause several interruptions every hour[1].

As shown, the performance requirements on the STL in a UDP type environment are very stringent. We can use either forward error correction (FEC) or automatic repeat requests (ARQ) to effective reduce the required bit error rate. FECs add redundant information into the data stream that allow the receiver to reconstruct the original message with corrupted or missing message parts. This technique adds additional bandwidth requirements to the data stream regardless of the actual error rate. Forward error correction works well on the physical link level where many bit errors can be present, but only provides diminishing returns in environments of lower bit error rates.

ARQ, on the other hand, relies on error detection at the receiver and issues a retransmission request back. This only requires additional bandwidth in the case of detected bit errors. This scheme, does require a feedback path in order to issue retransmission requests.

Even a single retransmission of a bit in error can greatly improve the effective bit error rate, as shown in the following example that provides a much better effective bit rate based on a channel bit rate of $7.65*10^{-6}$.

$$BER_{eff} = BER^2 = (7.65 * 10^{-6})^2 = 5.85 * 10^{-11}$$

Regardless of the type of error correction that is chosen, the impact of redundant data transmission or retransmissions can impact the exciter's synchronization ability, if it impacts clock packet transmission.

E2X TRANSPORT REQUIREMENTS

Overall the 3rd generation radio broadcast system architecture is required to propel digital radio to the next level. It is architecturally sound and the E2X protocol in itself works reasonably well. The protocol is not directly responsible for the previously illustrated challenges. Rather, the 3rd generation radio broadcast architecture is confronted with the challenges of high performance real time data streaming, a non-trivial challenge confronting the networking community. Good networking practices can ease the presented challenges, however, the problem is that there is no established transport protocol that adequately addresses all the requirements of the E2X protocol.

The E2X protocol poses the following transport requirements:

- Some packet types of the E2X protocol such as the initial control packet that carries critical system information, such as the configured service mode, require **guaranteed delivery reliability** in order to start digital modulation.
- E2X data packets only require **limited delivery reliability**, since data packets only convey meaningful information prior to the time that they

are intended to be on air. Retransmission requests beyond that time are futile and waste valuable bandwidth resources. A late delivery is essentially a packet discard.

- It is undesirable to retransmit any E2X clock packets, since any retransmission introduces a throughput delay measurement error. Depending on the implementation of the clock recovery circuit, most implementations can handle lost clock pulses better than incorrect clock pulses. **No delivery reliability** is, therefore, desired for clock packets.
- The protocol requires a **very high degree of reliability** in order to maximize the mean time between HD dropouts. If the transport protocol cannot provide an adequate degree of reliability, then the physical link layer is responsible for reliable data delivery instead.
- Clock packet delivery must not be impacted by
  - data packet transmissions, which can introduce a patterned throughput delay error that is difficult to recover from.
  - data packet retransmissions, which could make the throughput delay dependent on current link characteristics, such as environmental conditions.
  - aggressor traffic that is trying to share the link with the E2X protocol and may upset synchronization and even starve the digital modulator.
  - other E2X traffic sharing the link with the E2X protocol.
- **Unidirectional and bidirectional STLs** must be supported.
- **Multiple destinations** should be addressable to support multi exciter configurations, as well as, multi and single frequency networks.

A NEW E2X TRANSPORT PROTOCOL

Looking at this list of requirements it is apparent that neither TCP nor UDP adequately cover all items on the list. Nautel fully supports the existing E2X protocol design including the standard transport protocol choices. In addition Nautel has implemented a new transport layer protocol that addresses the E2X protocol transport requirements and allows multiple exciters to synchronize to the same Exporter synchronization stream while providing reliable data transfers to each destination.

The proposed E2X transport protocol implementation is layered on top of UDP, since UDP allows broadcast and multicast communication, however, it may be possible to port the protocol to the real time protocol (RTP) in the future. The protocol is best illustrated by looking at a main/standby exciter configuration across a generic STL as shown in Illustration 4.

The protocol is broken into a control plane and data plane. The data plane carries the encapsulated and segmented E2X protocol to both exciters. The Exporter can directly communicate with a single exciter using unicast IP, which directly maps a single IP address, such as 10.10.10.100 to a unique MAC address, such as 00:50:C2:59:70:8C. In order for the Exporter to use this mode of communications, a bidirectional STL is required to be able to use the address resolution protocol. Alternatively, the intermediate router can be programmed with a static MAC entry in the case of a unidirectional STL. To address multiple exciters simultaneously, the Exporter must resort to either broadcast mode (without an intermediate router), directed broadcast (through the router), or multicast IP to address only a subset of exciters.

The control plane communication from the Exporter to the exciters follows the same communication modes as the data plane, but the Exporter may choose to only directly address one of several exciters. For example, the Exporter may grant the ability to issue retransmission request only to a subset of listening exciters allowing for a highly scalable point to multipoint deployment of this protocol. For the time being this function is reserved but not implemented.

The control plane communication from the exciters back to the Exporter is always unicast directly to the Exporter. In the case of a unidirectional STL, this communication stream will not make it back to the

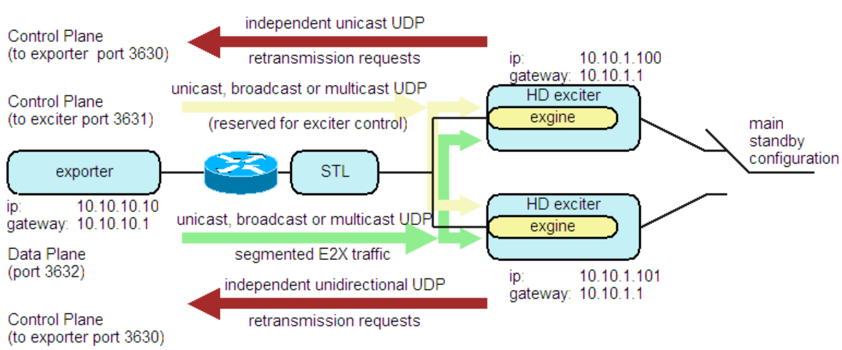

Illustration 4: E2X Transport Protocol for Main Standby Exciter Operation

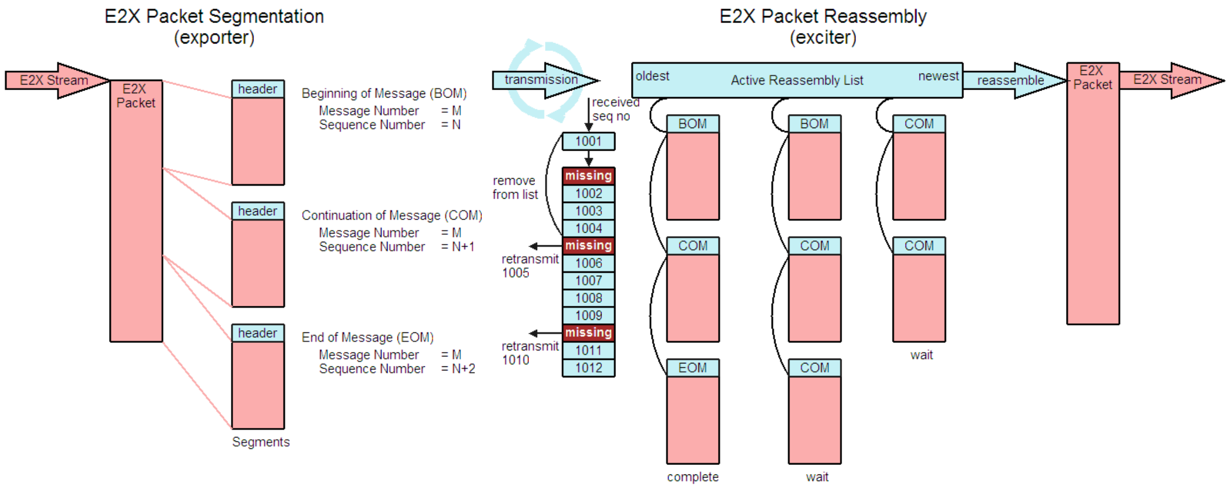

Illustration 5: E2X Packet Segmentation and Reassembly

Exporter, so the transport protocol cannot make the assumption that retransmission requests will in fact reach the Exporter. This will result in a loss of on-air quality due to increased data outages, but the E2X protocol continues to operate correctly.

Some packet types, such as the Exporter's control packet, require guaranteed delivery in order for the exciter to start modulation with correct operating parameters. Allowing exciters to start using locally stored operating parameters can cause a configuration mismatch between Exporter and exciter. Rather than relying on stored parameters, the E2X transport protocol maintains a copy of the most recent control packet and periodically retransmits the control packet to the exciter.

Using this configuration, both HD exciters are always ready to transmit HD and the digital modulator is not aware whether it is active or in hot-standby mode. The Exporter also is not aware which exciter is currently active. Both exciters independently issue retransmission requests back to the Exporter. If only a single exciter has lost a packet, it will issue a corresponding retransmission request. If both exciters lose a packet, both will issue a retransmission request but the Exporter will only retransmit the same data packet once provided the retransmission requests have been received close enough in time. Retransmitted packets are delivered to both exciters just like a regular data transmission. An exciter knows based on sequence numbers whether it has already received a packet and will discard duplicate data packets. Retransmission requests are made continuously until the data packet arrives or until the operation times out accounting for the fact that retransmission requests themselves may be lost.

This approach not only allows for redundancy in the exciters, it also allows the entire data path between Exporter and exciters to be split and made redundant. However, there is no guarantee that both digital modulators are always in the same state; IBOC sequence numbers and interleaving may be different and frame alignment may be offset. An exciter change-over normally causes an RF power cycle, so the exciter switch is not continuous in the first place. Receivers will detect the loss of digital coverage, blend back to analog and eventually re-acquire the digital signal.

E2X Segmentation and Reassembly

The main standby example demonstrates how the E2X transport protocol addresses multiple destinations and how it can support unidirectional and bidirectional links, it does not illustrate how the remaining requirements are met. The mechanics of the data transfer is shown in Illustration 5. In a real time data stream control over the PDU's payload size is crucial in order to guarantee a reasonable quality of service. Hence, the arbitrary length E2X packets are broken up into smaller segments with their own header information. Sequencing information in the header allows the segment to be related back to the original E2X packet.

The transmission path across the link may discard individual segments or reorder segments. Any segment that is received is queued in the active reassembly list in accordance with it's message number carried in the segment header. Once every segment in a reassembly list is sequential and the list consists of at least one beginning of message (BOM) and end of message (EOM) segment, the reassembly process is complete. The reassembled E2X packet is passed along to the digital modulator for processing.

In the event that the E2X packet to be segmented does not carry enough payload information to span more than one segment, the E2X packet is packaged into a single segment message, which then bypasses reassembly at the exciter and directly gets turned into an E2X packet. Clock packets fall into this category for faster processing.

Every sequence number that is received at the exciter gets entered into a sequence number list in ascending numerical order. After entering the sequence number,

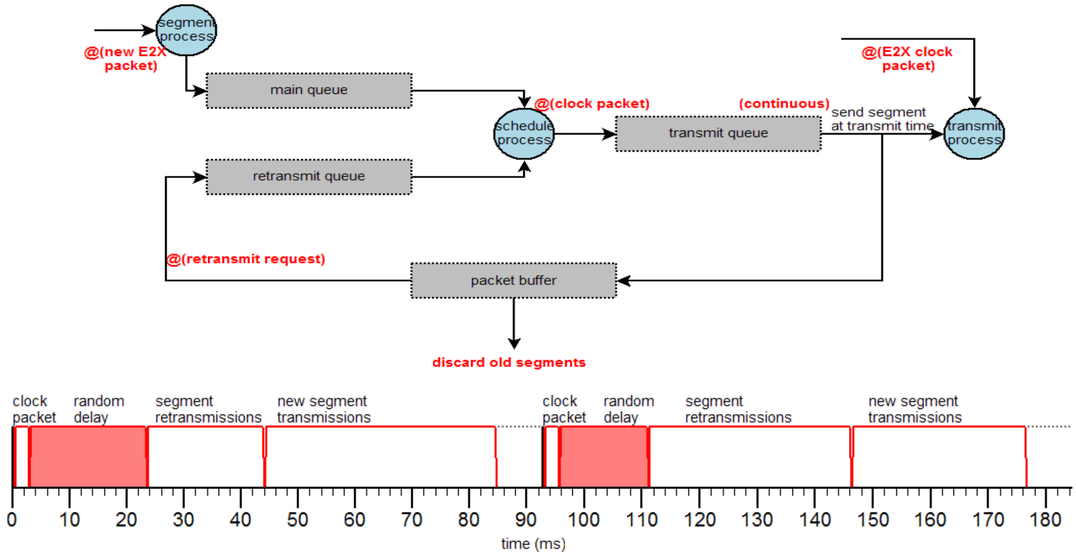

Illustration 6: Transport Protocol Scheduling at the Studio

the list is collapsed up to the next missing sequence number or to the end of the list. Now this list can be iterated across to determine what sequence numbers to retransmit. Multiple sequence numbers can be placed into a single retransmission request in order to minimize the number of retransmission requests.

While it is beneficial to decrease segment sizes in order to be able to control traffic flow more accurately, it does increase computational complexity, as well as, decrease bandwidth efficiency due to the added header structure. Larger segment sizes have a greater impact on the synchronization on another E2X stream on the same link. The segment size is configurable and not limited to Ethernet frame sizes; the underlying UDP protocol will break the segment into smaller parts automatically.

Segment Scheduling

Segmenting the E2X data stream in and of itself does not provide any benefits, since UDP does precisely perform the same function already. What complicates the situation is the fact that most STLs operate at a much lower network bandwidth compared to most LAN deployments. Consequently, all segmented packets will likely be placed on FIFO queues effectively negating the effect of segmentation unless the STL provides more sophisticated queuing. Illustration 6, however, presents a queuing structure that can effectively utilize the segments and compute packet dispatch times that match the reduced STL throughput rate. On the Exporter every new E2X packet is segmented right away and all resulting segments are placed on the main queue and each segment is treated independently from thereon.

The reception of a clock packet triggers the scheduling process that is controlled by two parameters:

1. The main bandwidth parameter sets the bandwidth required to sustain the regular E2X data stream without retransmissions. It determines how many segments may be transmitted in one clock interval and causes any instantaneous bandwidth variations to smooth out over time.
2. Overall STL bandwidth dedicated to this E2X stream dictates the rate at which segments are dispatched from the Exporter. It also governs how much additional bandwidth is available for retransmissions.

The scheduling process places segments from both the main queue and the retransmit queue on to the transmit queue as shown in Illustration 6. However, clock packets bypass the transmit queue and are directly sent to the exciter with a flag set to disabled retransmissions for clock packets. As each segment is queued on the transmit queue a dispatch time stamp is attached to it. The transmit process only dispatches segments after the given time stamp has elapsed.

The scheduler may first insert a random delay where no segments are scheduled. This minimizes the impact of one E2X stream onto another on the same STL, as it allows the two streams to interlock better while minimizing the synchronization impact. The scheduler first determines how much bandwidth is used up by new segments, but it schedules retransmissions prior to new segments, since retransmissions are more time critical.

All dispatched segments are placed on the packet buffer until they get too old to be of interest (i.e. greater than the exciter receiver buffer depth). A retransmission request from any destination searches the packet buffer for the given sequence number and places it on the retransmit queue. If two or more requests for the same sequence number from different exciters are received, only the first request is serviced, since the same same packet is delivered to both destinations.

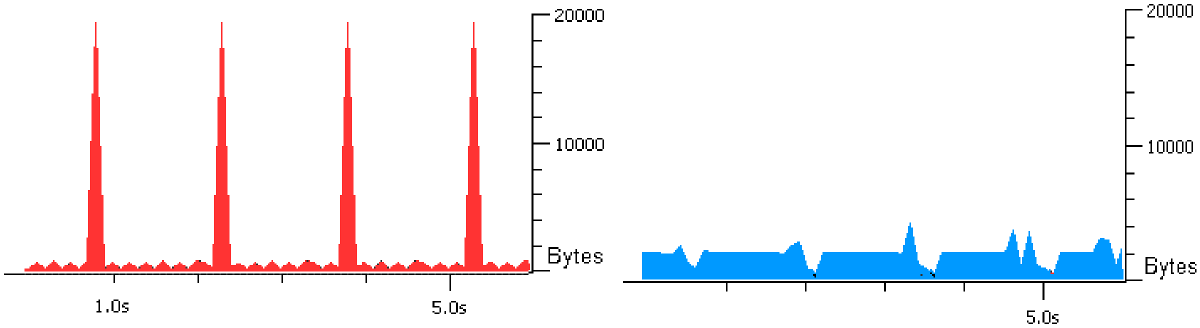

Illustration 7: E2X Instantaneous Bandwidth (left) and E2X Transport Protocol Bandwidth (right)

Illustration 7 shows a marked improvement in bandwidth utilization compared to standard E2X bandwidth utilization and demonstrates how multiple E2X streams can coexist on the same STL much more easily.

Bit and Burst Error Performance

In order to evaluate the effectiveness of this retransmission scheme, random bit errors are intentionally introduced at the segment level rather than causing discards. That way bandwidth resources are still consumed while not conveying any useful information. Illustration 8 contrasts the transport protocol's bit error performance with the E2X protocol in UDP mode based on the relationships discussed earlier. The exciter has a configurable receive buffer depth that allows the station operator to choose between increased protocol performance and total signal propagation delay. Two buffer levels of 16 packets (1.48 seconds) and 25 packets (2.32 seconds) are illustrated. Keep in mind that the E2X stream is spread almost over an entire L1 frame depending on how much bandwidth is allocated for it, reducing the effective buffer depth to a degree. This means that not every segment has the same chance of successful retransmission. With a buffer depth of 16 packets some packets may only be retransmitted once or twice; a buffer depth of 25 packets provides increased reliability due to a greater chance of successful retransmission.

The number of retransmissions depends on the throughput delay across the STL. The presented data takes this delay into account as it has been collected using a Moseley Starlink / Lanlink combination as presented in the application examples section.

A station engineer should analyze the effective bit error performance of his STL and select the desired quality of service. This identifies the operating point on this chart. If necessary the exciter's buffer depth can be adjusted. High quality RF/T1 based STLs do exist with bit error rates that are sufficient for good quality IBOC transmission. However, the STL itself is not the only source of data loss. Intermediate network equipment can contribute to packet loss due to congestion conditions or other circumstances; most network equipment is not designed to provide extremely high reliability, as in normal network situations, the end nodes provide reliable data delivery as part of their

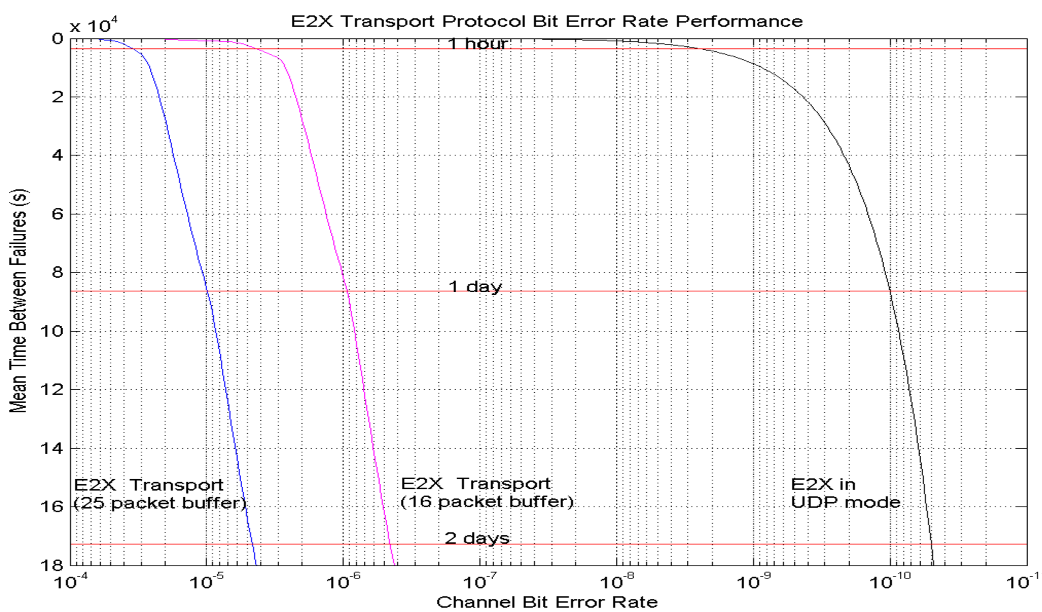

Illustration 8: Bit Error Performance

transport protocol. These symptoms will manifest themselves as packet loss which can be lumped into an effective bit error rate when averaging over a long enough time period.

Bit errors are a good way of characterizing the effectiveness of this protocol, but uniform random bit errors are almost an ideal error distribution, if they happen infrequently enough. In reality, bit errors are not uniformly distributed and usually occur due to an interference condition or fading. So it makes sense to investigate the maximum amount of time the exciter can maintain a steady data stream and replenish the data due to a complete STL interruption and associated packet loss. The following table outlines the burst tolerance of the protocol with respect to a configured receive buffer depth.

Buffer Depth (packets)	**Buffer Depth**	**Maximum Error Burst**	**Max Aggressor Traffic (300 kbps)**
16	1.48 s	200 ms	7.3 kB
25	2.32 s	600 ms	22.0 kB
35	3.2 s	1300 ms	47.6 kB
50	4.64 s	2100 ms	76.9 kB
75	6.96 s	3700 ms	135.5 kB

The leading reason for an STL interruption may not be the RF path at all. Instead the STL may experience a congestion condition due to other traffic, which could lead to packet loss or severe packet delay with the same consequences. The last column shows how much aggressor traffic can cause a corresponding STL interruption given an effective throughput rate of 300 kbps. The speed differential between common 10/100 Mbps LAN installations and common STL bandwidths can easily allow these amounts of traffic to be injected into the STL if no precautions are taken.

Quality of Service (QoS) Considerations

In order to effectively deal with aggressor traffic and share the available STL bandwidth with other traffic, absolute priority must be given to the E2X data stream. The IP protocol provides the Type of Service (ToS) octet to differentiate network packets and assign them different priorities. The IP ToS octet has the following definition:

0	1	2	3	4	5	6	7
precedence			delay	throughput	reliability	reserved	

The precedence field identifies whether this packet is routine traffic, immediate, or critical traffic. The delay, throughput, and reliability settings may either be set to high or low. However, these settings only have an effect if the intermediate network equipment provides the appropriate class of service queuing. Some types of STLs, such as the DMD20 satellite modem from Radyne provides a queue manager unit in its Ethernet bridge option that prioritizes traffic flow across four output queues using the IP4 ToS field to differentiate packets. Ideally more STL manufacturers would include QoS support in their STLs.

With QoS in place we now have a solid transport solution that allows the E2X stream to share the STL bandwidth with other E2X streams or even unrelated traffic, such as transmitter remote control.

APPLICATION EXAMPLES

With this information let's look at two application examples that use the E2X transport protocol.

Dual Station Topology

This topology is ideal for multiple HD stations that are co-located at the same studio and transmitter site. It demonstrates how multiple STL paths can be combined on the same network and is shown in Illustration 9. Two links with sufficient forward error correction to allow a reasonable quality of service, such as the Moseley Starlink, are chosen to carry the HD traffic to the transmitter site. These links have enough bandwidth to carry two HD streams simultaneously. Therefore, each link can serve as a backup STL to the other link. Because the E2X transport protocol is structured in such a way as to minimize impact on the other data stream, the presence or non-presence of the other data stream only affects the average clock packet throughput delay by about 4 ms; an amount that can easily be adjusted for to match the analog diversity delay. In conjunction with this protocol and a dedicated link, the StarLink provides clock packet jitter under 1 ms, this allows the exciter to synchronize to an error under 0.1 ppm. Therefore, no GPS locked synchronizer is required at the transmitter site to attain IBOC level II synchronization.

The network is broken into two distinct subnets the studio subnet and the transmitter subnet. This is necessary in order to isolate the three parallel paths between the two sites. Ethernet topology intrinsically does not like to be configured in closed loop systems as it could lead to unnecessary packet forwarding. Spanning tree protocols implemented in many Ethernet devices will intentionally sever Ethernet links if they detect multiple paths between two locations. By placing layer 3 routers between the two subnets not only are the subnets isolated, but we can also prevent any unnecessary network traffic, such as broadcast traffic, from traversing the link.

The forward links are unidirectional by nature, but the exciters for both stations can share the same bidirectional Moseley LanLink STL for retransmission requests. A host at the studio can select which path to take to the transmitter site by pointing its gateway setting to the appropriate router connected to the desired STL path. Only the Exporter points its gateway to the StarLink routers, the Importer or other

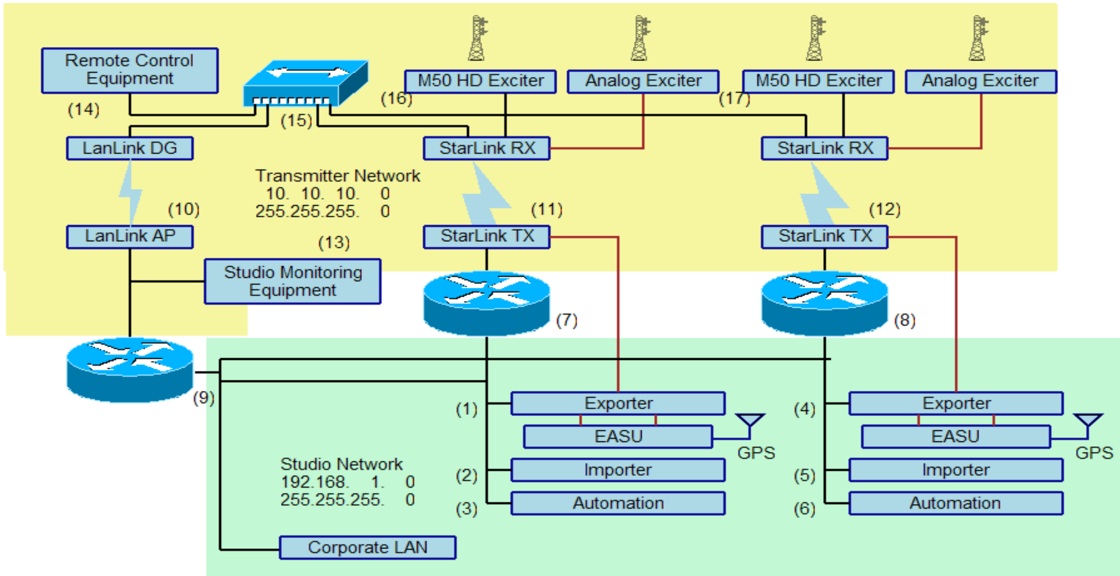

Illustration 9: Dual Station Topology

studio equipment would use the LanLink router to access the exciter or other control equipment.

The exciters, on the other hand, point to the LanLink router as their gateway in order to reach the Exporter for retransmission requests or serve control and status information.

This topology has many advantages:

- Redundant HD forward path allows the Exporter to select which link to use by simply reconfiguring its gateway. It can be demonstrated that this switch can be made quickly enough that with retransmissions enabled no on-air data loss occurs allowing STLs to be serviced and maintained.
- The reverse path can be shared by multiple stations reducing infrastructure costs.
- The LanLink STL is not on the critical HD path allowing it to be shared with other network applications without the use of quality of service queuing.
- Unlike TCP/IP, the E2X transport protocol is not dependent on the reverse path. For TCP/IP a failure (including congestion) in the reverse path would terminate the forward connection, as well. The E2X transport protocol loses the ability to retransmit data while continuing operation.
- Using directed broadcast, each station can have a standby HD exciter, provided that the studio side switch can prevent the broadcast stream from traversing the bidirectional LanLink STL back to the studio.
- Link monitoring can be included into the E2X transport protocol with support for automatic link switch over on link failure.

Satellite Distribution Network

Satellite modems are a convenient way of extending IP networks to remote locations, such as transmitter sites. Many satellite modems are bidirectional, such as the DMD20 satellite modem by Radyne. Even with the increased throughput delay, the proposed E2X transport protocol is well suited for satellite distribution.

Illustration 10 shows a proposed satellite distribution topology based on DMD20 Ethernet Bridge Option White Paper by Radyne[2]. In order to protect the satellite bandwidth, each satellite modem is connected to a router, effectively placing the studio and each transmitter site on its own subnet. This prevents broadcast or directed broadcast from being used, since broadcast traffic is only active within a given subnet. Broadcast E2X distribution could still be used by removing the routers at the transmitter sites, but using IP multicast is the preferred solution. The Ethernet standard provides for point-to-multipoint communications by assigning a special class of MAC addresses for multicast communications.

The IP multicast standard makes use of this fact and dedicates a class of IP addresses for multicast use that directly map to multicast MAC addresses. In this example, the IP address 232.0.0.1 is used, as it is a generic address that does not have to be registered with the Internet Assigned Numbers Authority (IANA). Multicast communications and multicast routing is fundamentally different from standard unicast communications, the multicast IP address does not identify the host, but rather identifies a data stream. So the assigned multicast IP address does not have to match the subnet IP address.

Different multicast routing protocols exist to pass multicast traffic from router to router. However, for our application it is sufficient to configure static multicast routes, for example "ip igmp static-group 232.0.0.1

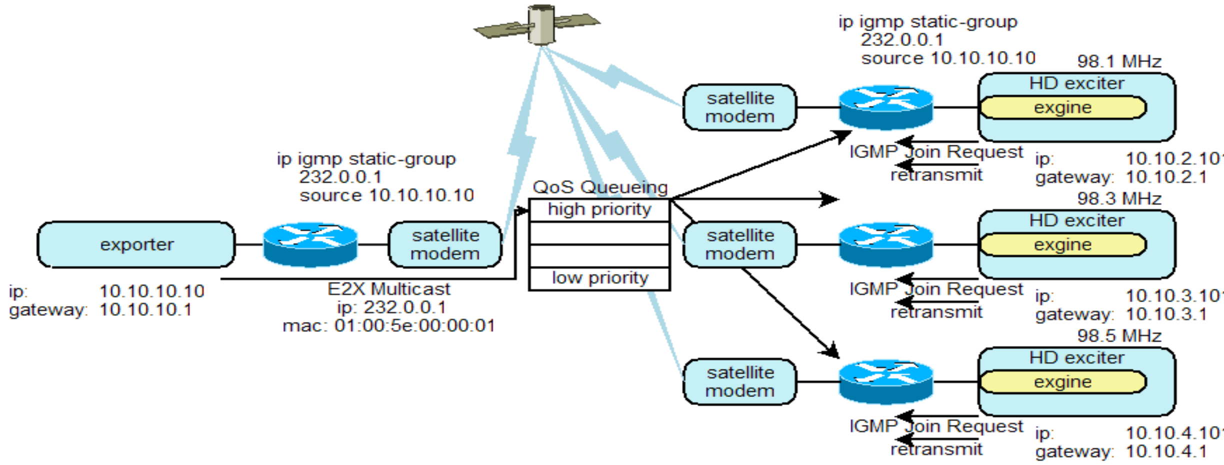

Illustration 10: Multi-frequency Satellite Distribution Network

source 10.10.10.10" configures a static route for Cisco routers.

If an exciter wants to join the E2X group, it issues a join request to its gateway using the Internet Group Management Protocol (IGMP). If the request is granted, the router will forward the multicast traffic to the exciter subnet. Only a single multicast data stream is transferred via satellite. Since the satellite modems employ quality of service queuing, we configure the E2X transport protocol to send clock packets at the highest priority and data segments at the next highest priority. All other traffic should be routed using a lower priority.

Each exciter will request data retransmissions using unicast communications back to the Exporter. So the intermediate routers must configure their appropriate routes back to the Exporter. All of these routes may also be defined statically.

In order to compensate for the increased throughput delays of satellite communications, the exciter buffer depth may need to be adjusted accordingly. To guarantee N=2 retransmissions per segment, which should be sufficient considering the high quality of satellite communications, we need to take the worst case round trip to and from the satellite into account. One way delays to a geostationary earth orbit (GEO) satellite is normally budgeted around 125 ms[3]. So for a successful retransmission, we require two trips to the satellite and two trips back to earth. Note that the initial throughput delay of the original packet is irrelevant. A single retransmission may take 500ms, adding 250ms of processing delays, leads to 750ms per retransmission. In order to ensure retransmissions across an entire L1 IBOC frame, we should add 1.48s. A buffer depth of 2.98s or 32 packets should be sufficient to ensure reliable retransmissions.

This proposed E2X transport protocol provides the reliable data transfer infrastructure to allow the deployment of multi frequency radio networks with each transmitter broadcasting the same content at different frequencies. Single frequency networks may provide an effective way of extending digital coverage in the future, but before they can become a reality, the IBOC standard must define L1 absolute frame alignment, ensure exciter synchronization, and define signal delay characteristics. It is clear, however, that single frequency networks will require a robust studio to transmitter data path, since the loss of a single data packet on a particular path could cause the digital modulators to assume different state potentially requiring the entire network to be restarted.

Single frequency networks will likely require a GPS reference at each transmitter site, but depending on the variance of the satellite throughput delays, multi frequency networks may not require a GPS reference.

CONCLUSION

The proposed E2X transport protocol provides drastic improvements over current E2X deployments, reducing bandwidth and bit error rate requirements on the studio transmitter link while allowing for Ethernet based synchronization. It provides additional flexibility in the layout of STLs by allowing multiple streams to be co-located on the same STL. By implementing a feedback path back to the studio automatic repeat requests can be used to greatly improve quality of service.

By incorporating this protocol into the radio system broadcast architecture, various multiple exciter configurations can be implemented, such as main standby exciters, multi frequency networks and single frequency networks.

REFERENCES

1 iBiquity Digital Corporation, October 2006
HD Radio™ Data Network Requirements
http://www.ibiquity.com/hd_radio/iboc_white_papers

2 Radyne ComStream, Inc, March 2004
DMD20 Ethernet Bridge Option

3 Freeman, R. L., 1997 Radio System Design for Telecommunications 2nd ed. New York: John Wiley & Sons