The Exciter

The etymology of the word *exciter* may have been lost to time, although the answer is probably found in early FM broadcast equipment manuals. Its context to the broadcaster has been in the generation of an FM signal at low power, apparently to *excite* the FM transmitter into oscillation. The Radio Amateur's Handbook of 1946[1] leaves this impression:

> *"A complete transmitter therefore may consist of an oscillator followed by one or more buffer amplifiers, frequency multipliers, or straight amplifiers, the number being determined by the output frequency and power in relation to the oscillator frequency and power. The last amplifier is called the final amplifier, and the stages up to the last comprise the exciter."*

Fast-forward a couple of generations, and the term "exciter" is equated with a box that generates an analog FM radio signal at low power. It could drive a vacuum tube resonant amplifier or a broadband solid-state amplifier; or it could simply have a harmonic filter placed on its output to be a low power transmitter. Thus, the meaning of the word has evolved. Adding a new device to the category, with the second generation (GEN 2) of HD Radio transmission devices, the IBOC digital signal generator became known as an exciter. Just like an analog exciter, the GEN 2 exciter takes an audio input and puts out an IBOC digital radio signal.

The Exgine

At an FM station converting to GEN 2 HD Radio operation, the existing analog exciter was paired with a new GEN 2 HD Radio exciter, and their two independently generated signals were combined in one of the ways discussed in the previous chapter. As the products evolved, the analog exciter and digital generator could occupy the same device—the exciter would take in two audio inputs and output the hybrid signal. Then the third generation of HD Radio transmission equipment was made possible when iBiquity separated the front end from the back end of the HD Radio Exciter. On the back end, the end that generates the radio signal, transmitter manufacturers developed analog FM exciters that could accept an optional HD Radio generator board operating at the lowest layer of the IBOC protocol, designed simply to create the IBOC OFDM waveform. To provide the input to the generator, a separate unit was developed that handled the front end—processing the higher layers of the IBOC protocol, starting with the audio inputs. Considering the back-end OFDM generator be an IBOC *engine* that works inside an analog exciter, iBiquity formed a contraction of the two words, calling this IBOC building block the exciter engine, or Exgine (pronounced EX-jin). The Exgine generates the digital signal, passing on to the amplifiers in the exciter. The Exgine-equipped exciter can transmit a digital-only signal or a hybrid digital signal. (It can transmit analog-only, too, but there would be no need for the Exgine card to do that.)

Connecting the Services to the Transmitter

In this section on transmission, the focus is on the input to the Exgine and the exciter. iBiquity has a fine Application note on the subject of IBOC system input topologies, so we won't go into excruciating detail here. Look for *HD Radio Data Network Requirements,* iBiquity document number

---

[1] The Radio Amateur's Handbook, 1946 Edition, American Radio Relay League, p. 94, From the library of Lewis D. Collins, Sc.D.

TX_TN_2040, at www. iBiquity.com. Before reading that Application Note, the reader should be familiar with the service modes (such as MP1) and the services (such as MPS or SPS), which are threshed out in the chapters of this book.

Our focus will be on FM IBOC topology. Providing input to AM IBOC systems is a relatively uncomplicated affair compared to FM IBOC systems. AM IBOC involves transmitting only main program audio and data services (MPS and MPSD), plus the station information service (SIS). In addition to MPS, MPSD, and SIS, FM IBOC systems can transmit the optional supplemental services and advanced data services. To generate these additional services requires extra building blocks in the broadcast facility.

The simplest configuration of an IBOC input system is identical to that of an analog system. Main program audio and program-associated data are delivered to the transmitter site on the STL and are plugged into the back of the IBOC exciter. Figure 14.1 illustrates how the existing transmission infrastructure of an analog station can be utilized to transmit the host analog signal along with the IBOC MPS, MPSD, and SIS. The studio audio program is transmitted to the transmitter site in the usual way, via studio-to-transmitter link (STL). The low bit rate data for the MPSD is typically multiplexed on the STL as well, or delivered by other means. SIS data can be preprogrammed into the transmitter, requiring no STL bandwidth. Thus, with the simple addition of a slow data interface, the studio side of the system is ready for IBOC.

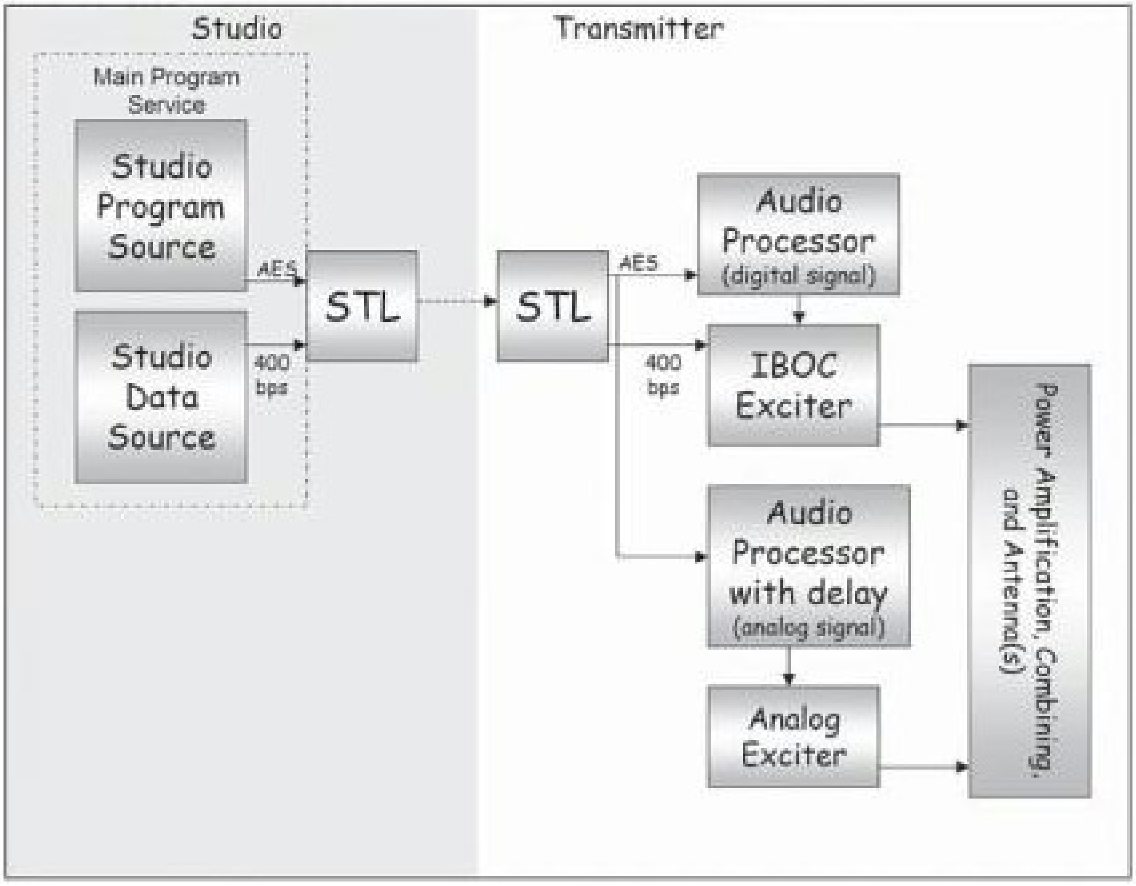

Figure 14.1 Simple IBOC Transmission Topology with All IBOC Gear at Transmitter Site

Main Program Audio

At the transmitter site, the data goes into a port in the IBOC exciter. The audio, however, is subjected to several processes before transmission. It must be split into two audio streams, one each for analog transmission and digital transmission. Of utmost importance, the analog audio must be delayed to accommodate the latency of the digital audio and to provide a diversity delay to overcome brief catastrophic reception failures. The audio clocks at studio and transmitter are best synchronized by frequency locking their clocks (often to a GPS-derived source) or processing the digital audio stream through a synchronized rate converter to avoid buffer underflows or overflows on the air. (Audio may also be sent across an analog STL, obviating studio-transmitter clock issues by having all digitization occur at the transmitter site; this comes with the disadvantages of an analog audio plant and STL.)

Also very important, each audio stream requires an audio processor to tailor it to the transmission medium and the "sound" of the station. Processing for a digital transmission is very different than processing for an analog system. Several factors play into audio processing for analog FM audio, including the prevalence of FM noise and multipath characteristics, as well as the tension between the high frequency pre-emphasis and the 100% modulation limit. On the digital side, the codec provides a predictable and repeatable noise and distortion environment, generally unaffected by the over-the-air medium.[2] As long as the data stream is decodable, the full audio bandwidth and full signal-to-noise ratio of the audio is available to the receiver. This provides the digital audio stream with a different set of constraints on its audio processing than those that apply to the analog audio.

Analog and digital audio signals are subject to maximum levels in different ways. On the digital side, it is a numerical limit on the maximum output level of the analog-to-digital converter—the converter counts perfectly to the maximum level and can go no higher. On the analog side, audio amplifiers in the transmission chain will more gracefully overload than a digital converter does, first producing compression effects, then clipping. Also, over the air, analog transmission has a regulatory limit of 100% modulation, requiring the analog audio processor to maintain compliance by compression and/or clipping. For all these reasons, separate audio processors should be applied to the digital and the analog audio streams.

---

[2] Hybrid AM IBOC has two digital audio streams, a basic core stream and a lower reliability enhancement stream. The audio quality will change in the absence of the enhanced stream. Nevertheless, the signal-to-noise ratio of even the core audio is consistent and predictable.

Economizing for Multicasting

With the advent of multicasting, the demands placed on the STL increased substantially. If full-bandwidth audio streams were required for the supplemental audio channels, broadcasters would have to double up on their STL equipment and bandwidths. Why not take advantage of the bit-rate reduction of the HDC coder at the studio? Analog audio for the main program would still need to be delivered to the transmitter site, but with no such requirement for supplemental services, the concept made sense as a way to minimize the bandwidth required of the STL.

To support SPS channel coding at the studio, the protocols for generating an HD Radio signal would have to be divided between two locations—the studio and transmitter sites. Doing so could also benefit the broadcaster who wanted to control the main program digital and analog audio streams at the studio. By encoding the MPS audio at the studio, the station's audio processors and time alignment functions could remain at the studio. The extra bandwidth required for the MPS digital stream over the STL is about 120 kBps.

A distinction is made here between the NRSC-5 protocol stack and the HD Radio protocol stack. The NRSC-5 stack begins by receiving encoded audio from outside NRSC-5 and placing it on the audio transport, within NRSC-5. HD Radio transmission devices start by encoding the audio with the HDC coder, then passing it to the audio transport. As we talk about the input stages of the IBOC system, the HD Radio protocol stack includes audio encoding processes not included in NRSC-5.

The Exporter

To separate protocol functions between studio and transmitter, a new device has to be installed at the studio, and some of the functions normally performed at the transmitter must be stripped and moved to the studio. Clearly, the HDC coder must now be resident at the studio. It also makes sense to, at the studio, combine the program service data with the audio, making it ready for transport to the transmitter site and eventually over the air. The *HD Radio Data Network Requirements* document describes the management of data flow from studio to transmitter by logical channel (P1, P3, PIDS). Recall from Chapter 5 that these logical channels are formed by the Layer 2 multiplexing process. Therefore, it can be inferred that some or all of Layer 2 processing occurs at the studio side of this split system.

In Figure 14.2, the division of the HD Radio protocol stack between studio and transmitter sites is illustrated. A new device, the Exporter, performs the HD Radio processing tasks for the MPS, MPSD and SIS. The Exporter concept and software is an iBiquity creation that transmitter manufacturers employ in their HD Radio products. The Exporter performs several functions. It passes the incoming main program audio feed to the digital audio processor and the analog audio processor. The diversity delay is applied to analog audio. Digital outputs include the audio stream for analog transmission (not coded) and an Ethernet feed of the IBOC digital logical channels. These can be conveyed to the transmitter site by conventional private STL methods or by commercial digital telecommunications links. Alternatively the Exporter can be installed at the transmitter site as if it were part of the transmitter.

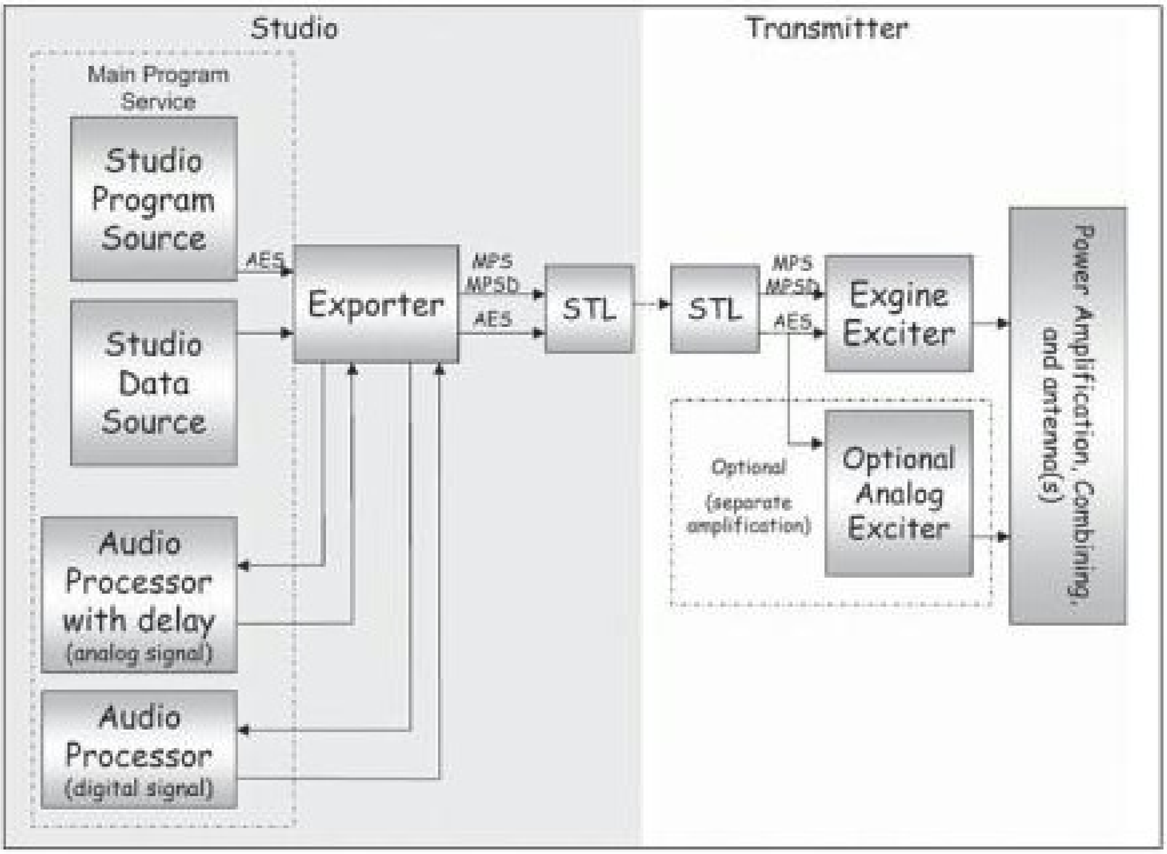

Figure 14.2 IBOC Transmission Topology with HD Radio Exporter, and Exgine Modules

Not shown is the optional use of a GPS receiver to drive the studio AES/EBU audio clock at the same rate as the transmitter's. (AES/EBU is the familiar name for a series of digital audio signaling formats jointly adopted by the Audio Engineering Society and the European Broadcasting Union.) In the functional block of the Exporter in Figure 14.3, the Exporter is depicted as a single block containing several processes. One of the processes depicted in the figure is not an original Exporter function. It is the function of the Exciter Auxiliary Service Unit (EASU), also called a *Synchronizer*. The EASU is available as a separate device, but can be integrated into one package with the Exporter. Functionally, the main program audio goes into the EASU/Synchronizer stage for clock synchronization, and goes out for audio processing. The block marked "sync" illustrates that. Then the EAS/EBU streams are processed by the Exporter functions. The EASU is GPS-synchronized to provide the both the AES/EBU digital audio stream and the Exporter a clock that matches the clock at the transmitter. The Exporter inputs labeled "From Importer" are discussed in the next section titled Importer.

the inference that Layer 2 functions are incorporated in the Exporter. On the other hand, Layer 1 convolutional coding of P1 increases the bandwidth by the 2/5 coding rate, to about 246 kbps, more than the required average bandwidth would allow on the STL. We can infer that Layer 1, with its convolutional coding, is executed at the transmitter.

This splitting of IBOC protocols between two sites presents two critical challenges—timing and bit-accuracy. If the information does not make its way to the transmitter in a consistent and timely manner, the IBOC transmission will hiccup. There will be a discontinuity in the audio. If the information is inaccurate and compromised by errors along the way, packets may be rejected or possibly sent with errors, resulting in irregular performance as well.

A New Transport Where None Had Been Necessary Before

To split elements of the IBOC protocol stack (in this case, Layers 1 and 2) between two geographic locations requires yet another transport protocol. Initially, to move information in a timely manner, employing a one-way STL, the User Datagram Protocol (UDP) was chosen to transfer the information across the STL. UDP is one of the Internet protocols,[5] and might be more aptly indicated as UDP/IP, the way the more familiar TCP/IP protocol pair often is. UDP is a low-overhead protocol that simply embeds a checksum in its packets for error detection and it requires no response from the recipient. It is useful in situations where speed is valued over accuracy, on one-way links, and in a network broadcast where the overhead of responses from multiple recipients would be burdensome. Early adopters of the UDP transmission for their IBOC signals quickly learned how important it is for the STL to have as high reliability as possible.

HD Radio Data Network Requirements spells out the reliability issues. Running a 48 kbps coded audio stream via UDP, "any dropped IP packet will result in the loss of 1.5 seconds of HD Radio programming." A 0.1% packet loss would result in 873 dropouts per day. A 0.001% packet loss reduces that to 9 per day, and 0.0001% to about one per day.

On the other hand, with a bidirectional STL, and a little more bandwidth, the TCP protocol employs a feedback mechanism to retry lost or failed packets. So long as there is enough time to resend missing packets, TCP provides a much more robust link. *HD Radio Data Network Requirements* says as much as 1% packet loss can be accommodated on a TCP STL, with a 20-frame buffer and an average stream rate less than 60% of maximum. 1% loss is regarded as an "unhealthy LAN." Less than 0.1% loss is regarded as healthy. TCP is demonstrably the more reliable way to transmit IBOC frames from studio to transmitter, but it, too, has its limitations. For instance, by the time a discarded packet is retransmitted, it may be too late, wasting bandwidth on information that cannot be utilized.

---

[5] Request for Comments 768, Internet Engineering Task Force, tools.ietf.org/html/rfc768, August 1980.

It is less critical when the STL is dedicated to the HD Radio traffic and on its own reserved subnet, in which case even a 10 Mbps network is "more than plentiful."

Those who have observed the output of the Exporter have noted the high burst of packets that comes once a frame period. Initially there was some concern that the STL would require enough capacity to replicate the bandwidth of the burst. However, the only reason for the burst lies in the manner in which IBOC Layer 2 works. Recall from Chapter 5 that it accumulates an entire frame of data to pass along to Layer 1 as a monolithic PDU. Once the PDU is ready, it is shoved out the door at the maximum rate the Exporter and its network interface can handle. MTM looked into this phenomenon to see how to ensure broadcast network configurations were compatible with it. They conclude that the burst bandwidth concern is unfounded. Networking devices are designed with bandwidth disparities in mind. When a faster link (the Exporter to a switch or bridge) is connected to a slower link (the STL), MTM explains that Ethernet specifications (IEEE 802.1)[7] require the slower link to be able to buffer 2 seconds of input. The burst from the Exporter is smoothed out by the buffer in the bridging device (the switch) and allowed to trickle through to the STL at its rate. Since each frame is less than 2 seconds duration, and the STL has the bandwidth to handle the required average data rate, the buffered data will make it across the STL without congestion. This conclusion presumes that all traffic on the STL has been accounted for in selecting the STL bandwidth, and that the bridging device is actually compliant with the 802.1 buffering criterion.

Network Security

MTM also discusses network security in *HD Radio Networking Best Practices*. While its study of several HD Radio facility networks was not a security audit, MTM did notice that security on the radio stations' networks was "practically nonexistent." At best stations had assigned passwords to administrative features, but they were often the same password on all parts of the network. MTM recommends among other things, diverse passwords across the network, maintaining most up-to-date patches and bug fixes for any network code in use, centralized authentication, authorization and accounting (AAA) to control and audit network activity, encryption of network management traffic, limit administrative access between subnets to only those paths that are necessary, assign read-only privileges where appropriate.

Another Approach

In a paper presented at the 2007 NAB Broadcast Engineering Conference, Philipp Schmid of Nautel Limited, revealed his analysis of the Exporter transport weaknesses.[8] Nautel developed an E2X Transport Protocol that utilizes UDP to move data and a new protocol as a transport to wrap around the UDP. It employs two UDP forward streams—one for timing and control and one for

---

[7] IEEE Standard for Information technology- Telecommunications and information exchange between systems- Local and metropolitan area networks. Part 5: Remote Media Access Control (MAC) bridging; IEEE Std. 802.1G, Institute of Electrical and Electronics Engineers, standards.ieee.org.

[8] E2X Bandwidth and Bit Error Requirements for Ethernet Synchronization; Introducing a Reliable Real-Time Point-to-Multipoint E2X Transport Protocol, *NAB Broadcast Engineering Conference Proceedings*, Philipp Schmid, 2007, National Association of Broadcasters.