

Optimizing Analog + HD Radio Transport Over an Existing 950 MHz STL Channel

Bob Band, Junius Kim & Keyur Parikh

Harris Corporation
Mason, Ohio

INTRODUCTION

This paper addresses the challenges involved in adding HD Radio transport over an existing 950 MHz STL path. Beyond simply adding an IP stream alongside the analog FM, we'll look at what can be done if we re-examine the whole concept from the ground up with joint FM/HD transport in mind. What should the next generation of 950 MHz STL look like? What combination of technologies, both new and established, provide the best recipe for success?

We will look at such issues as: How much traffic can a 950 MHz STL support? What backhaul alternatives might be used to enable TCP traffic without a full 950 MHz return path - and how can this provide network-based remote management for monitoring and control? What techniques allow us to carry other Ethernet (LAN) traffic alongside the HD stream without interference between them? What are the tradeoffs between modulation waveforms, data payload, and required channel bandwidth? Equally importantly, we'll look at ways to make this all work on existing licensed spectrum, using existing cables and antennas.

THE EXISTING 950 MHz STL

FCC Part 74, Subpart E authorizes the frequency band of 944-952 MHz for use by radio stations for over-the-air transport of audio program material, with the most common application being the STL, or Studio-to-Transmitter Link. The carrier frequency channel spacing is 25 kHz, with common licensed bandwidths being 200, 300, and 500 kHz.

Before the advent of the digital STL, analog 950 MHz STLs were used for transport of the baseband FM composite signal. The digital STL offers significant advantages in terms of flexibility, reliability, and quality. It can multiplex many user channels, such as one or more high quality audio programs, Radio Broadcast Data System (RBDS) data, and Ethernet traffic over the same 950 MHz RF link, and being digital, this multiplexing is seamless – no crosstalk or degradation between adjacent channels. The digital STL also offers the possibility

of lossless data transmission; that is, errorless end-to-end data transport. This is in contrast to analog STLs where if some degradation of the RF signal occurs, this loss is transferred to the demodulated FM composite signal.

This digital multiplexing capability means that in theory, any digital 950 MHz STL system should be capable of carrying the Ethernet data involved in getting HD Radio on the air.

In practice, however, there are a number of areas in which problems can arise, so in order for an STL to be truly “HD ready” it behooves us to look at what these areas are and how we can best overcome the difficulties

THE HD RADIO SIGNAL

First of all, let's define what we mean by the HD Radio signal to be carried on the STL. This can actually be any of several different signals, depending on the placement of the HD Radio Importer and Exporter.

There are three basic scenarios:

1. If the Importer and Exporter are both placed at the transmitter site, then the STL needs to carry one or more audio programs (one for the FM/HD1, and optionally those for HD2 and HD3) plus a low-speed data signal for the Program-Specific Data (PSD) and, if implemented, Advanced Application Services (AAS) traffic for data-only services.
2. If the Importer is at the studio and the Exporter at the transmitter site, then we need to transport the Importer-to-Exporter (I2E) signal, which is a TCP data signal that can run as high as 156 kbps. Being TCP, it requires duplex connectivity.
3. If both the Importer and Exporter are placed at the studio, we need to carry the Exporter-to-Exgine (E2X) signal. This can be either UDP (one-way) or TCP (duplex), and its bandwidth varies depending on the Service Mode in use (MP1, MP2, MP3, or MP11).

In the worst case scenario, MP11 using TCP, we need about 330 kbps of duplex bandwidth for E2X transport. This latter approach is the one recommended (though not required) by iBiquity.

What all of these have in common is the requirement for IP data transport along with audio - and, in at least some cases, duplex IP data transport.

To properly design a new approach to the STL, we need to look at two distinct but interdependent areas of technology: digital RF transmission and data communications.

DATA BANDWIDTH REQUIREMENTS

As we have just seen, depending on the HD component placement, the STL will need to transport up to 330 kbps of IP data traffic along with one or more program audio channels. For the main analog FM broadcast program, an uncompressed linear stereo pair with 16-bit samples and 32 kHz sample rate (typical for a 15 kHz FM signal) requires just over 1 Mbps of bandwidth; increase the sample rate to 48 kHz for 20 kHz audio and the requirement jumps to 1.5 Mbps. Well-established compression algorithms such as those in the MPEG and apt-X families can reduce these requirements considerably, but it is always wise to keep the STL uncompressed wherever possible in order to minimize the chances of compromising the audio quality. Add to this the transport of auxiliary audio channels such as Subsidiary Communications Authority (SCA) and RBDS, and we see that the total data bandwidth requirements can easily exceed 2 Mbps.

CARRYING CAPACITY OF THE STL

Three basic factors affect the amount of data we can carry across the STL link: the amount of RF

bandwidth available, the modulation scheme in use, and the RF quality of the link.

The first of these is simply the amount of RF bandwidth you've requested (or been able to get) under your Part 74 license. For FM stations this is usually 300 kHz, but can be as little as 200 kHz or as much as 500 kHz. The theoretical maximum amount of data that can be carried in each of these depends on the type of modulation chosen for the link.

Modulation

Quadrature amplitude modulation (QAM) is a modulation scheme widely used in STLs and many other communication applications. In digital telecommunications, QAM data is sent as binary symbols in a grid with the number of points in the grid being a power of 2 (2, 4, 8, etc). Common QAM forms are 16-QAM, 64-QAM, 128-QAM and 256-QAM. By moving to a higher-order symbols or constellations, it is possible to transmit more bits in the same RF bandwidth. However, in doing so, the mean energy between points in the QAM constellation must be closer together and thus more susceptible to noise and other corruption. This results in a higher bit error rate (BER) so that all other things being equal, a higher-order QAM can deliver more data than lower-order QAM – but it does so with less reliability.

Additionally, in a practical QAM application, periodic pilot symbols are inserted in the payload. These pilots carry no user data and are used for enhanced phase noise immunity and improved adaptation or acquisition by the receiver. However, careful design and optimization of the RF receiver can reduce the overhead used for pilot symbols.

Table 1 shows the spectral efficiency and maximum data carrying capacity under ideal conditions for various QAM orders.

QAM Order	Spectral Efficiency, Bits per Second per Hz	Raw Data Bandwidth (Mbps) with a 200 kHz RF bandwidth	Raw Data Bandwidth (Mbps) with a 300 kHz RF bandwidth	Raw Data Bandwidth (Mbps) with a 500 kHz RF bandwidth
4	2	0.4	0.6	1
16	4	0.8	1.2	2
32	5	1	1.5	2.5
64	6	1.2	1.8	3
128	7	1.4	2.1	3.5
256	8	1.6	2.4	4

Table 1: QAM and Spectral Efficiency

IMPROVING THE RF LINK QUALITY

For a given RF path, a study is usually done to determine the total required RF system gain which factors in parameters such as antenna gain, cable loss, transmitter power, over-the-air distance, fade margin, required BER etc. The RF STL hardware plays a critical role in this system gain equation.

Design and optimization of a digital RF STL is a tradeoff between spectral efficiency and BER. The difference in required signal to noise ratio when moving between different QAM orders is around 3 to 4 dB per step. So, for example, 128-QAM will require 3 to 4 dB more signal to achieve the same BER performance as 64-QAM.

There are several ways to increase the performance of a given RF link. These include:

- Using a more powerful RF transmitter, and/or higher gain antennas
- Optimizing the RF hardware circuits to minimize any inherent circuit noise
- Adding some form of forward error correction (FEC)

The first two of these are common to any RF system and present no special challenge in the design of a combined FM/HD transport system. FEC, however, deserves a closer look.

Essentially, in FEC we send some additional information on the link along with the payload (audio and data) information, and use this information to detect data errors and rebuild any missing information at the receive end.

There are several types of error correction available, the most common and best known of which is Reed-Solomon. However, FEC adds overhead to the total data transport requirement, and thus introduces a Catch-22: if you're trying to compensate for poor link quality caused by using high QAM orders, and you're using those high QAM orders in order to carry more information, adding FEC can obviate the information rate gains. The key, therefore, is to use the most efficient form of FEC possible, with efficiency being defined as the ability to correct the maximum amount of errors while sending the minimal amount of FEC overhead.

One of the most efficient FEC schemes is low-density parity-check (LDPC) codes or turbo codes.

In information theory, the noisy-channel coding theorem (aka Shannon's theorem) establishes that however contaminated with noise interference a communication channel is, it is possible to

communicate digital data (information) nearly error-free up to a given maximum rate through the channel. The Shannon limit of a communications channel is the theoretical maximum information transfer rate of the channel for a particular noise level.

While LDPC and other error correcting codes cannot guarantee perfect transmission, the probability of lost information can be made vanishingly small. LDPC was the first code to allow data transmission rates close to the theoretical maximum, the Shannon limit. In fact, it has been shown that LDPC codes can reach within 0.0045 dB of the Shannon limit. Figure 1 shows the signal to noise ratio (SNR) vs. spectral efficiency for LDPC coding.

Although the powerful mathematical concepts behind LDPC have been understood for decades, it is only recently, with the advanced computing power available in today's digital signal processing chips, that practical implementations have become possible.

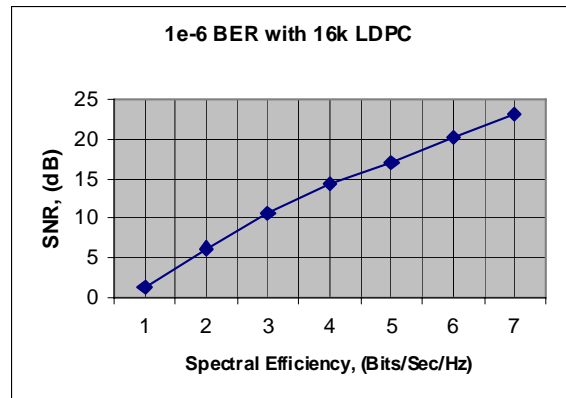


Figure 1: SNR for 1E-6 BER vs. Spectral Efficiency

In contrast to less efficient error correction schemes such as Reed-Solomon, LDPC can offer a clear advantage in terms of performance for a given signal to noise ratio. Specifically, for equal amounts of FEC overhead, LDPC requires 3dB to 5 dB less signal for the same BER performance than does Reed-Solomon. In the real world, this can translate into using the next-higher QAM order while maintaining the same BER and RF power level.

NETWORK LAYER

Now that we've seen the potential for a big improvement in spectral efficiency on the RF side of things, let's look at what might be done on the data side to make our new STL better suited to its more complex role.

The network layer capability of the next generation STL is one of its core components. This set of functionalities should not only meet the unique challenges that are posed by a typical STL transport

application, but also provide a reduction in capital and operational expense by integrating advanced data networking features.

To this end, the STL should contain an IP gateway architecture that supports bandwidth management to efficiently prioritize and transport critical and non-critical traffic across the STL link, and be able to operate in one-way as well as full duplex link configurations.

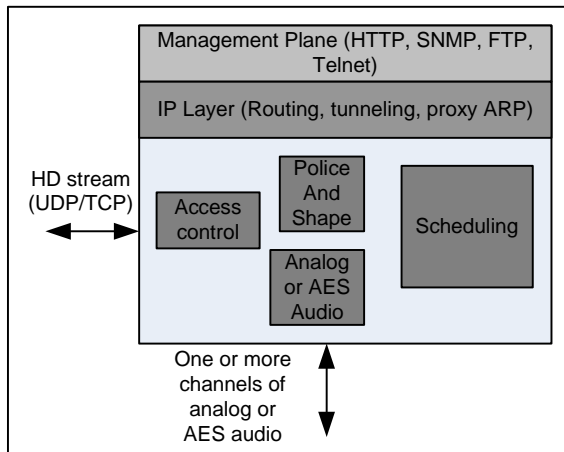


Figure 2: STL IP Gateway Architecture for the Next-Generation STL

Figure 2, above, depicts a high level view of the functionalities that could be designed into the next generation STL's IP gateway. In addition to the core features, the platform manageability of this gateway contributes directly to minimizing the operational cost of the device. By supporting standard network management protocol functions such as HTTP, SNMP, FTP, Telnet and ICMP, the user is better able to effectively troubleshoot and manage the device.

DATA BANDWIDTH MANAGEMENT

After the RF element design of the STL achieves the maximum spectral efficiency, the network layer must then be designed to move the various types of data across this link in the most efficient way.

The integrated IP gateway needs to manage the available bandwidth and allocate it appropriately among the following traffic types: Audio for the FM and, depending on the component locations, the HD program(s); IP data for the I2E or E2X HD Radio signal; and other IP data, such as control and management, and even LAN traffic.

As the traffic from different sources converge on the IP gateway's data plane, it is imperative that the gateway's architecture support functionalities that will filter unwanted Ethernet traffic, police and shape the incoming Ethernet traffic and finally prioritize

traffic so that the low priority traffic does not interfere with the media related traffic. To achieve this, a committed traffic rate must be allocated to the various Ethernet traffic sources, as well as to the digitized audio channels. The traffic from the audio source(s) are well controlled and at a constant rate, and therefore require no metering or policing function; however, the Ethernet traffic may be bursty or variable in rate, and needs to be handled carefully.

As each IP stream (whether HD media or control traffic) enters the gateway's Ethernet interface(s), it should first pass through an access control function which will only pass packets whose IP source and destination address pair are configured as valid connection endpoints. This prevents unwanted traffic from a local subnet, including broadcast packets [Note: The term "broadcast" is used here in the data networking sense, meaning packets that are addressed to multiple recipients, not in the more traditional sense of radio or television broadcasting] to enter the STL link and eat away at bandwidth that is allocated for the legitimate media to be transported. Once the Ethernet traffic has passed access control, it needs to be prioritized as either High or Low priority based on either the physical port it has entered through, 802.1p priority, or the IP DiffServ code point.

A policing and shaping function ensures two things: one, that any excess Ethernet traffic entering the system via a lower-priority port, such as general LAN traffic, is discarded, thus preventing such traffic from interfering with the transport of the critical audio and HD streams. Second, the smoothing function accommodates instantaneous bursts of the HD's media stream across the allocated air bandwidth. Without this smoothing function, the transmission of instantaneous bursts in the HD IP stream will require far greater air bandwidth. Finally, a scheduling function should ensure that the strict priority is followed when queuing packets for transport across the RF STL link.

All the above functionalities ensure that the STL's forward bandwidth is efficiently utilized and provides network layer protection for the media traffic against any unwanted LAN event.

NETWORK TOPOLOGY SUPPORT

STLs can be deployed in two flavors: simplex or duplex. For simplex configuration, only a forward link from the studio to the transmitter is provided, which is then used for transporting UDP-based HD media and the digitized analog or AES audio streams to the transmitter site. However, increasingly the HD media stream is employing TCP as a transport protocol to reduce packet loss, which then requires a full duplex STL.

Since the 950 MHz paths are generally assigned only in the forward STL direction, the reverse link (from transmitter to studio) can be achieved in any of several ways: it can employ radios in other bands, such as the license-exempt 900 MHz ISM band, which is close enough to the licensed 950 MHz band to “piggyback” on the same antennas and cabling; other license-exempt radio links that may be in use for LAN/WAN traffic; an inexpensive DSL link from a managed IP service provider, as shown in Figure 3 below; or via a corporate WAN.

Additionally, the forward direction of these alternative IP paths could be used as a backup STL in case of catastrophic loss of the 950 MHz link. An intelligent STL design would enable automatic switching to this backup path as needed – and not just of the HD IP stream, but of the digitized FM program audio as well.

In supporting these different network topologies, the IP gateway in the next-generation STL needs to integrate all the necessary networking features that facilitate smooth operation. For example, a proxy-ARP feature can be integrated, which would provide MAC addresses on behalf of the nodes that are on the other side of the STL, thus avoiding the cost of an additional networking node such as a router for this purpose.

On a full duplex link, a layer 3 switching function should be provided so that the TCP return traffic from the HD node at the transmitter site traverses the IP return path back to the studio as discussed above. This requires that the STL’s IP gateway support multiple network interfaces and be able to efficiently switch traffic based on IP addresses. When the return path traverses a service provider’s network, the IP gateway must ensure that the private IP addresses of the HD nodes are kept private. This can be

accomplished by having the IP gateway send the return TCP traffic from the transmitter site through a UDP tunnel using IP addresses that are assigned by the service provider.

The tunnel endpoint is the studio site IP gateway which de-encapsulates the tunneled packets and routes them appropriately to the HD node on the studio network (see Figure 3). The tunnel encapsulation ensures that the packets are routable and traversable across Network Address Translator (NAT) devices in the service provider’s network.

IN SUMMARY

In looking at how best to design the next-generation 950 MHz STL, we need to pay close attention to several factors:

- Maximizing the spectral efficiency of the RF link through use of efficient digital modulation schemes such as QAM, while implementing state-of-the-art, highly efficient FEC methods such as LDPC.
- Controlling and managing data traffic to ensure that the audio and HD streams always get the highest priority
- Integrating Layer 3 switching functionality to ensure optimum performance on any of the types of low-bandwidth return links available to broadcasters today.
- Taking maximum advantage of alternate IP paths to provide integral STL backup in case of the failure of the 950 MHz link.

A new 950 MHz STL, one designed from the ground up for transport of HD Radio signals along with FM, is long overdue.

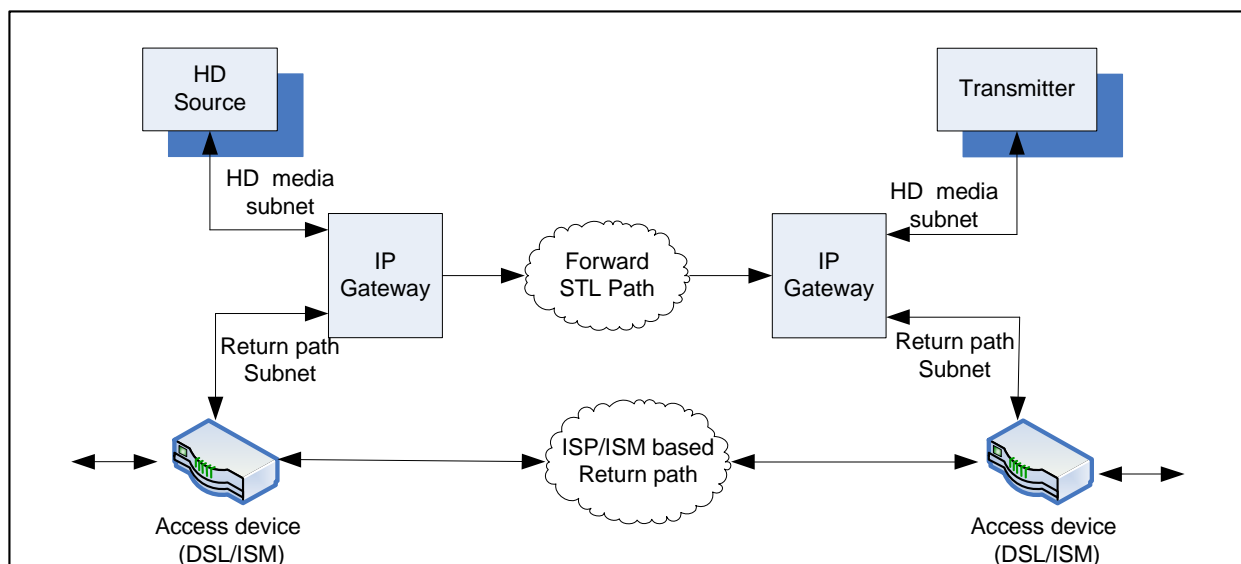


Figure 3: Network topology with return path