The ATSC Distributed Transmission System and Applications to Translator Service

David L. Hershberger
dhershberger@axcera.com
Axcera, LLC
960 McCourtney Road, Suite C
Grass Valley, California 95949-7423

Abstract

Distributed transmission is single frequency network technology applied to the ATSC system for digital television. Rather than using a single transmitter to service a coverage area, multiple transmitters are used. The transmitters are synchronized in frequency and symbol emission. Timing adjustments allow optimization of the system to produce minimum timing skew in areas where the multiple signals overlap.

This paper provides a detailed introduction to the technology and hardware developed for distributed transmission.

The first station to implement distributed transmission is WPSX-DT in State College, PA. For WPSX, distributed transmission was the only feasible way to provide UHF coverage, mainly because of terrain shielding. This paper presents implementation experiences and data from this real-world installation.

Distributed transmission can also be applied to translator systems, by creating distributed translator networks. In a distributed translator network, the translators may all operate on the same channel. This results in greater spectral efficiency. Multiple hop distributed translator systems can be accommodated with a minor change to the ATSC CS/110 candidate standard. The changes necessary to implement distributed translator systems are also described.

What It Is

Distributed Transmission is the use of multiple transmitters to service your coverage area rather than the traditional single transmitter system.

What It Is Not

Distributed Transmission is not a booster (repeater) system. Distributed Transmission systems cannot normally receive a signal off-air and retransmit it, because this would violate causality. Distributed transmission systems radiate the same symbol sequence at very nearly the same time from each transmitter. The relatively long time delays associated with demodulation, deinterleaving, and error correction, followed by forward error correction, interleaving, and remodulation, would mean that a distributed transmission slave would need the input data long before it could receive it over the air.

On channel repeaters are another beast entirely. On channel repeaters receive an off-air signal and retransmit it within about a microsecond. Most of the delay in an on channel booster is from the IF bandpass filter (usually a SAW filter).

Because the delay of an on channel booster is intended to be short, it is not possible to regenerate the data and remove errors. Received signal distortions (including the short echo produced by the booster itself) are cumulative and are simply retransmitted.

A distributed transmission system transmits a pristine signal from all of its transmitters. And, each transmitter can be delayed or ***advanced*** in time with respect to the other transmitters in the network. On channel repeaters can only be delayed.

Why It Is Needed

Some DTV stations are “terrain-challenged.” Most DTV allocations are on UHF channels, where terrain shielding is more of a problem than it is on VHF. In some cases, it is impossible to replicate analog service using a single transmitter. For these cases, distributed transmission is needed. In other cases, such as the inability to put up a single tall tower, distributed transmission may be desirable.

Putting up a distributed transmission system is more complicated than just feeding the same SMPTE 310 signal to multiple transmitters. In fact, two exciters from the same manufacturer, fed with the same SMPTE 310 input signal, will almost always generate different symbol sequences. In other words, they will become mutual jammers. DTV modulators are generally nondeterministic.

A DTV modulator makes an arbitrary decision on where to insert frame sync. Frame sync occurs once every 624 MPEG packets. So the odds of getting two modulators to insert frame sync in the same place is 1 in 624.

DTV modulators also include arbitrary initial states in 24 trellis coder bits and 12 precoder bits. (To simplify the terminology in this paper, we will collectively refer to these as “trellis coder” states, data, or bits, even though precoders are also included.) So, this makes a total of 36 arbitrary bits. So the odds of having two exciters fed with the same bit stream producing the same symbols by chance is 1 in $624*2^{36} = 1$ in 42,880,953,483,264, or about one in 43 trillion. Starting up an ATSC modulator once a second would result in chance synchronization about once every 1.4 million years. The odds get far worse when we consider having three or four modulators randomly synchronize.

Another way to say this is to point out that given a particular SMPTE 310 input stream, there are 42,880,953,483,264 different symbol sequences that can represent the same signal.

So to create a distributed transmission system, there must be a means of synchronizing the otherwise arbitrary initial conditions in each modulator, and maintaining synchronism over time. Pilot frequencies also need to be locked, and timing must be adjustable and stable.

How It Works

Figure 1 shows the architecture of a distributed transmission system. New hardware components are shaded. The new components include a distributed transmission adaptor (DTxA) and distributed transmission (DTx) slave transmitters.

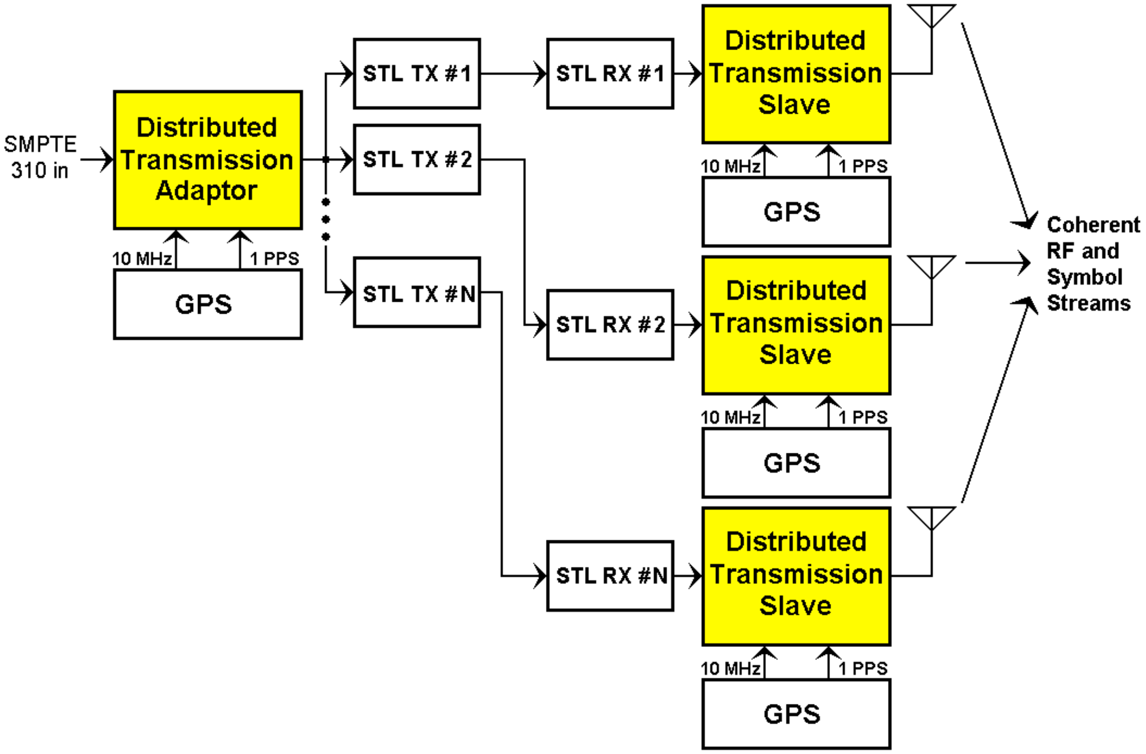

Figure 1 – Components of a Distributed Transmission System

A distributed transmission system also includes a GPS receiver at the system origination point and at every distributed transmitter site. The GPS receiver provides both frequency and time references. A 10 MHz frequency reference is used to lock up the IF signal and all of the local oscillators. The GPS receiver produces a 1 pulse per second (PPS) (1 Hertz) time reference, which is used as the timing reference for the network.

Distributed transmission uses a device called a distributed transmission adaptor (DTxA), shown in Figure 2. The input and output of the DTxA is SMPTE 310. The DTxA also requires the 10 MHz and 1 PPS GPS signals. The DTxA inserts data into a reserved packet in the SMPTE 310 bitstream that provides synchronization information to all of the transmitters in the network. This particular packet is called the distributed transmission packet or DTxP. A placeholder packet, including the correct packet identification (PID) word is inserted by the station’s service multiplexer. When the DTxA encounters this placeholder, it overwrites it with the rest of the data that forms a valid DTxP.

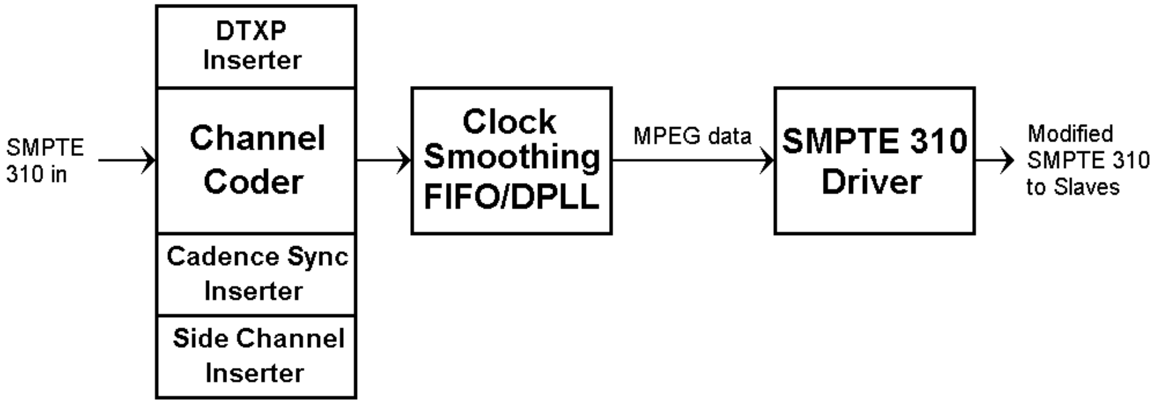

Figure 2 – Distributed Transmission Adaptor

The DTxP contains information that allows all of the transmitters to synchronize their trellis coders and frame sync insertion points. This ensures that all of the transmitters are emitting the same symbols.

The DTxP also contains timing information, with different settings individually addressed to every transmitter in the network. This allows the network to be tuned so that received timing skew is minimized in the most populated areas.

Another function of the DTxA is to smooth out frequency variations in the SMPTE 310 clock. A FIFO and a slow phase locked loop (PLL) perform this function.

The distributed transmitters in a network receive the modified SMPTE 310 signal from the DTxA. A block diagram of a “slave” modulator appears in Figure 3 below.

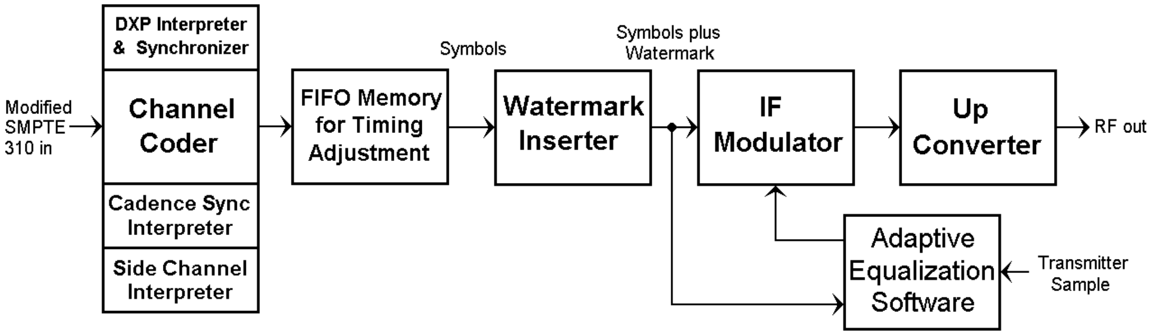

Figure 3 – Distributed Transmitter

The slave transmitter retrieves data from the DTxP and uses that data to synchronize its trellis coders, and to adjust its timing.

Another component of the distributed transmission system is cadence sync, which is separate from the DTxP. ATSC modulators normally make an arbitrary decision on where to insert their frame sync. The frame sync includes equalizer training sequences, VSB mode bits, and reserved bits. The frame sync is inserted once every 624 MPEG packets. (MPEG packets, in turn, each contain 188 bytes.) The cadence sync is a bitwise inverted MPEG sync byte at the point where frame sync is to be inserted. The MPEG sync byte, normally 47h, is changed to B8h.

RF Watermarking

There is another optional feature of a distributed transmission system. In a distributed transmission system, it is desirable to be able to identify signal components emanating from different transmitters. The basic purpose of distributed transmission – synchronizing frequencies and symbols – thwarts this purpose by making all of the transmitter output signals basically identical. So, how can you tell them apart when you are trying to adjust a network?

By using a watermark signal, it is possible to create a display on a test receiver that would look like the example shown in Figure 4. The contributions from each transmitter are identified by means of the RF watermark, which is different for each transmitter in a network.

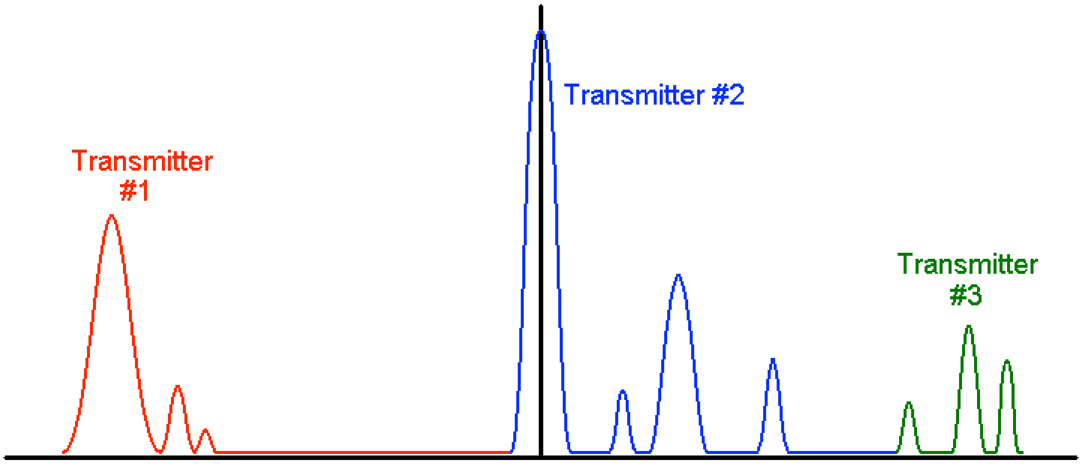

Figure 4 – Aggregate Channel Impulse Response with Contributions from Three Transmitters Identified

To aid in system setup and diagnostics, the distributed transmission system includes a unique "watermark" signal for each transmitter in a network. The watermark signal is a low-level code, buried underneath the ATSC symbols. It appears as noise to receivers. The noise level is small enough so that it will have little or no practical effect on threshold (typically only about 0.1 dB).

The RF watermark sampling instants are the same as the ATSC symbols. This creates eight minor eye openings in the demodulated I channel, with the normal seven major eye openings interspersed as shown in Figure 5. This figure was produced by a software simulation of the RF watermark injected at a bury ratio of 27 dB.

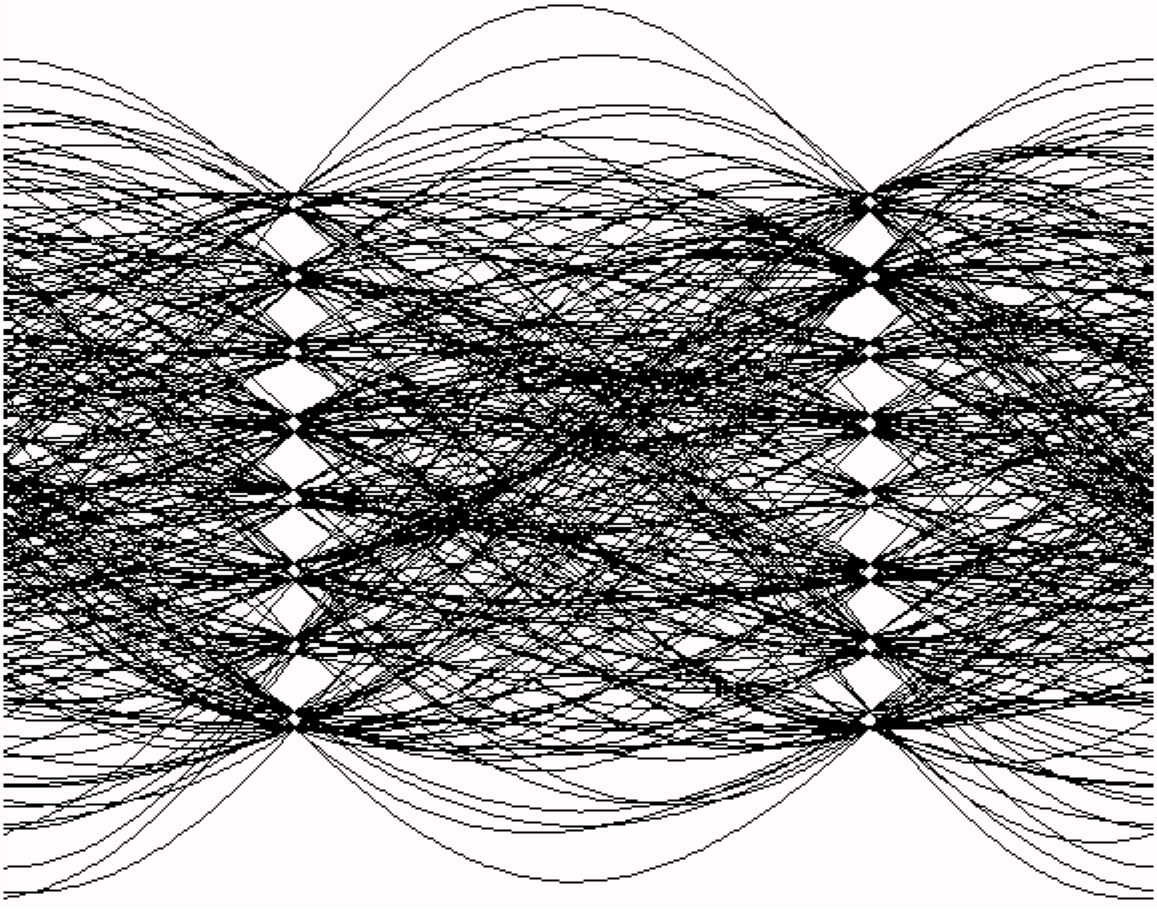

Figure 5 – Demodulated Eye Simulation with –27 dB RF Watermark

The RF watermark signal will typically be operated about 30 dB ***below*** the transmitter's average output. However, the coding benefits from a 54 dB coding gain when integrated over one ATSC field. So, it is theoretically possible to identify a buried code from a transmitter that is 24 dB weaker than an interfering transmitter in just one field. Additional averaging, over multiple fields, would increase the decoding gain.

A constellation display derived from the same simulation is shown in Figure 6. This display shows that each I channel line splits into two. The distance between the two constellation lines is proportional to the RF watermark injection level.

The objective in creating the RF watermark system is to be able to identify individual components of a received aggregate signal, and to associate the received components with individual transmitters, as shown in Figure 4 above. Even if one transmitter's signal is buried underneath a stronger signal, it is still possible to identify the components from the weaker transmitter because of the coding gain of the watermark signal.

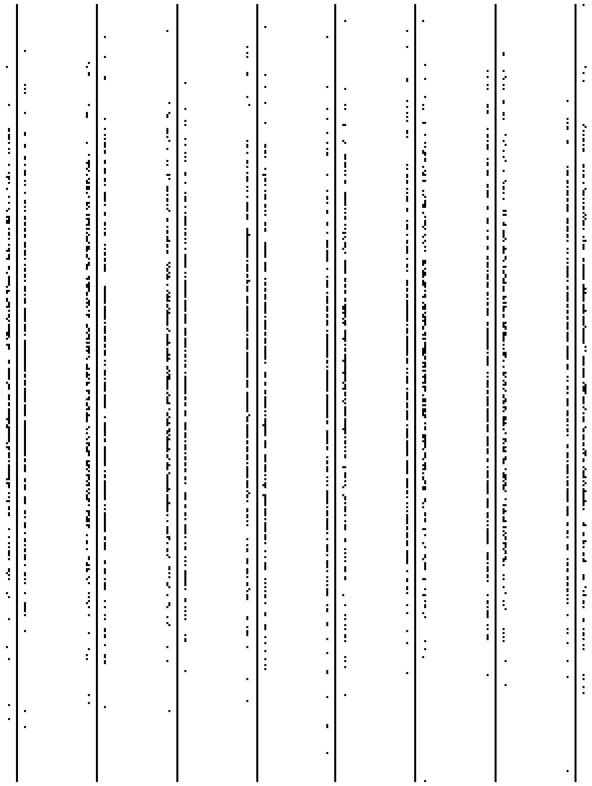

Figure 6 – Demodulated Constellation Simulation with –27 dB RF Watermark

Two things happen to the transmission system when an RF watermark is being transmitted. First, the RF watermark, since it looks like noise, establishes a signal to noise ratio (SNR) "measurement floor."

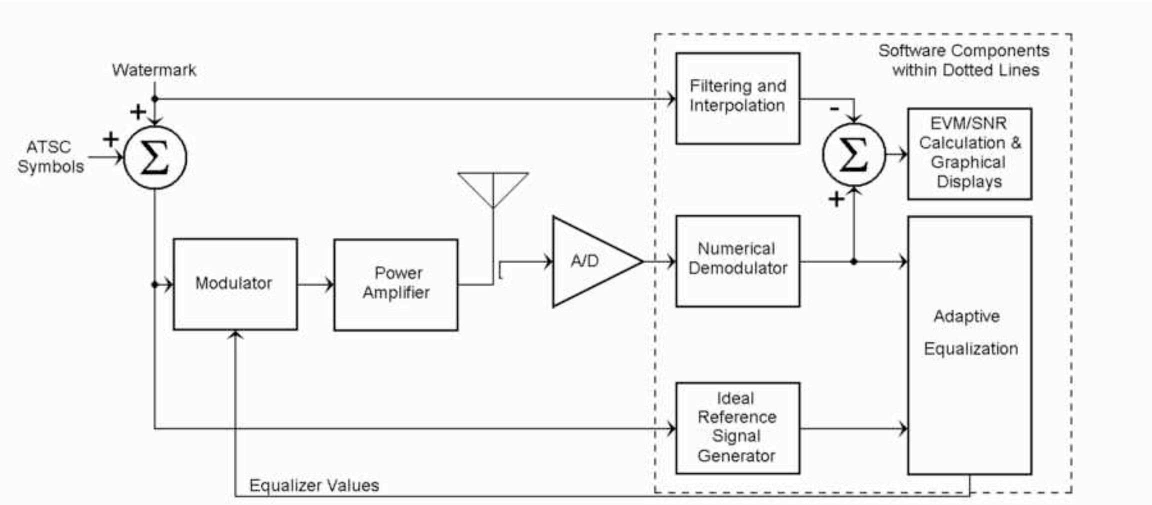

Figure 7 – Watermark Processing in a Distributed Transmitter

Although the measured SNR values that include the RF watermark are relevant to receiver performance, they will not accurately reflect transmitter performance. (Transmitter performance will be better than the measured SNR values.)

Second, when an RF watermark is being transmitted, particularly at the higher injection levels, it can have an effect on adaptive linear and nonlinear equalization systems. Fortunately, the transmitter knows what the watermark code sequence is, which means that the RF watermark itself may be subtracted from both the demodulated and ideal reference signals as shown in Figure 7. Subtraction of the RF watermark sequence not only results in improved adaptive equalization, but it also allows generation of SNR and Error Vector Magnitude (EVM) values, and constellation and eye diagrams without the watermark. With the watermark removed, the actual transmitter performance can be more accurately measured.

The distributed transmission packet (DXTP) is transmitted over the air, so everything in it is public information (with the exception of the trellis codes, which must be removed because of causality constraints). So, a test receiver can read the DXTP, find out what watermark codes are being transmitted, and search for them to determine which transmitters in a network are being received.

The level of the RF watermark may be set to one of seven levels, or it may be turned off completely. The different levels give the system operator flexibility to trade off watermark signal strength versus receiver noise threshold change. For example, when you are adjusting a network and want to be able to see weak transmitters, you might want to increase the amplitude of the watermark. Then during normal operation you may want to set it to a low level.

One more function of the DTxA is frequency control. There are several contributors to frequency inaccuracy in a digital transmitter. These include the IF modulator and the up converter. Figure 8 shows the system blocks that can affect frequency accuracy.

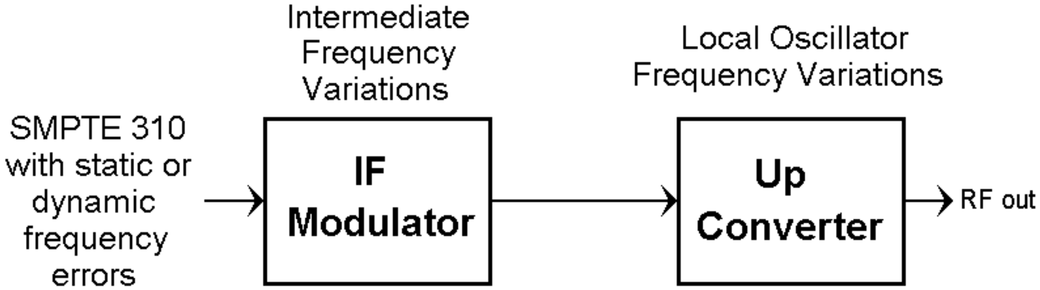

Figure 8 – Possible Contributors to Frequency Errors

Of these, the IF modulator is usually the main contributor to frequency errors. The up converter usually includes straightforward PLLs that can easily be referenced to an external GPS-derived 10 MHz clock. But the IF modulator may be affected, at least on a transient basis, by SMPTE 310 errors.

The frequency dependencies of an IF modulator are illustrated in Figure 9.

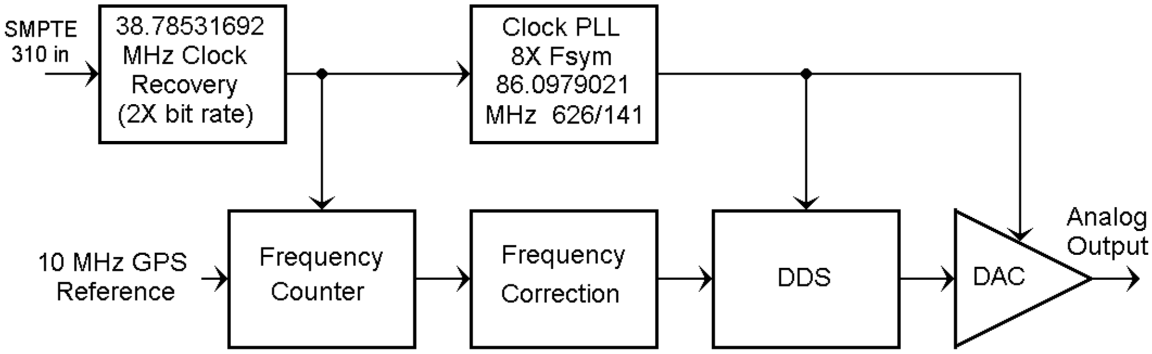

Figure 9 – How Modulator Clocks are Locked to SMPTE 310 Clock

This is not a diagram of an ATSC modulator, but rather a functional diagram that shows how the various frequencies in an ATSC modulator are related to the SMPTE 310 and 10 MHz reference inputs. Basically, all clocks are locked to the incoming SMPTE 310. The ATSC symbol rate is the SMPTE 310 clock frequency times 313/564.

In Figure 9, a clock recovery circuit locks a PLL to twice the SMPTE 310 clock frequency. Then a second PLL locks an ATSC symbol rate related clock to the SMPTE 310 clock. Eight times the ATSC symbol rate is the SMPTE 310 clock times 626/141.

When the input SMPTE 310 clock frequency changes, the IF spectrum of some ATSC modulators would move unless frequency correction is applied. An uncorrected spectral shift is shown in the Figure 10. The frequency shift is exaggerated.

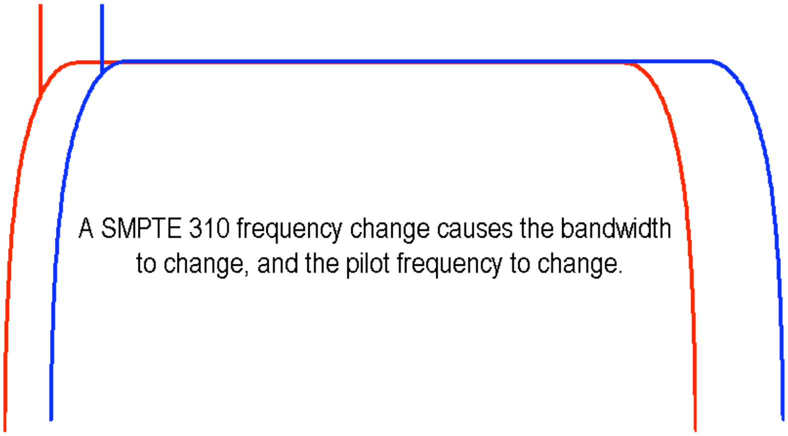

Figure 10 – Effect of Clock Frequency Changes on IF Frequency

Most ATSC modulators include frequency correction circuitry to make the IF frequency independent of the SMPTE 310 clock. However, there is a time lag associated with this approach. It takes a short time to measure the SMPTE 310 clock frequency, and then apply the necessary correction.

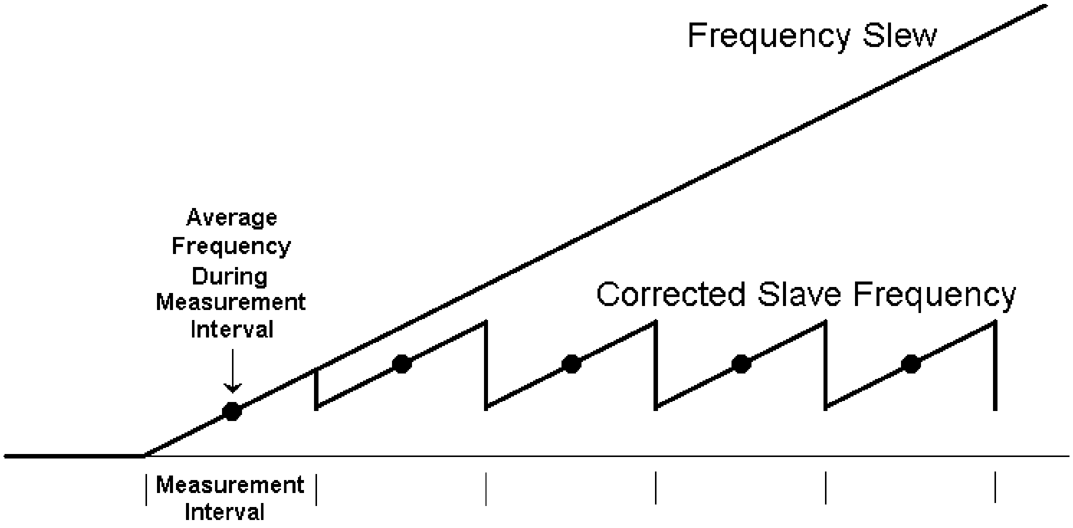

Figure 11 – Correcting Frequency Errors when SMPTE 310 Clock Frequency Is Slewing

As the frequency of the input SMPTE 310 slews, a simple frequency correction system may create a static frequency error. When the frequency error slewing stops, then the system will not create a static frequency error. The effect of frequency slew on a frequency correction system is shown in Figure 11.

One of the objectives of the DTx system, and particularly for the design of the DTxA, is to smooth the frequency error slew enough that the static frequency error shown above is sufficiently small that it will not create problems. That means being within 0.5 Hz of the nominal IF frequency.

This can be accomplished by using a very slow PLL, with a time constant of about 20 minutes. Such a slow PLL will keep the SMPTE 310 frequency slew within the limit of 0.028 parts per million (ppm)/second even if the input is hot switched from a –2.8 ppm to a +2.8 ppm source (the frequency error limits). The transient response of such a PLL is shown in Figure 12.

The horizontal scale in Figure 12 is one hour. Because this PLL is so slow, it is implemented in digital circuitry. Such a slow analog PLL would have required very large low leakage capacitors, and initial acquisition would have required large charging currents.

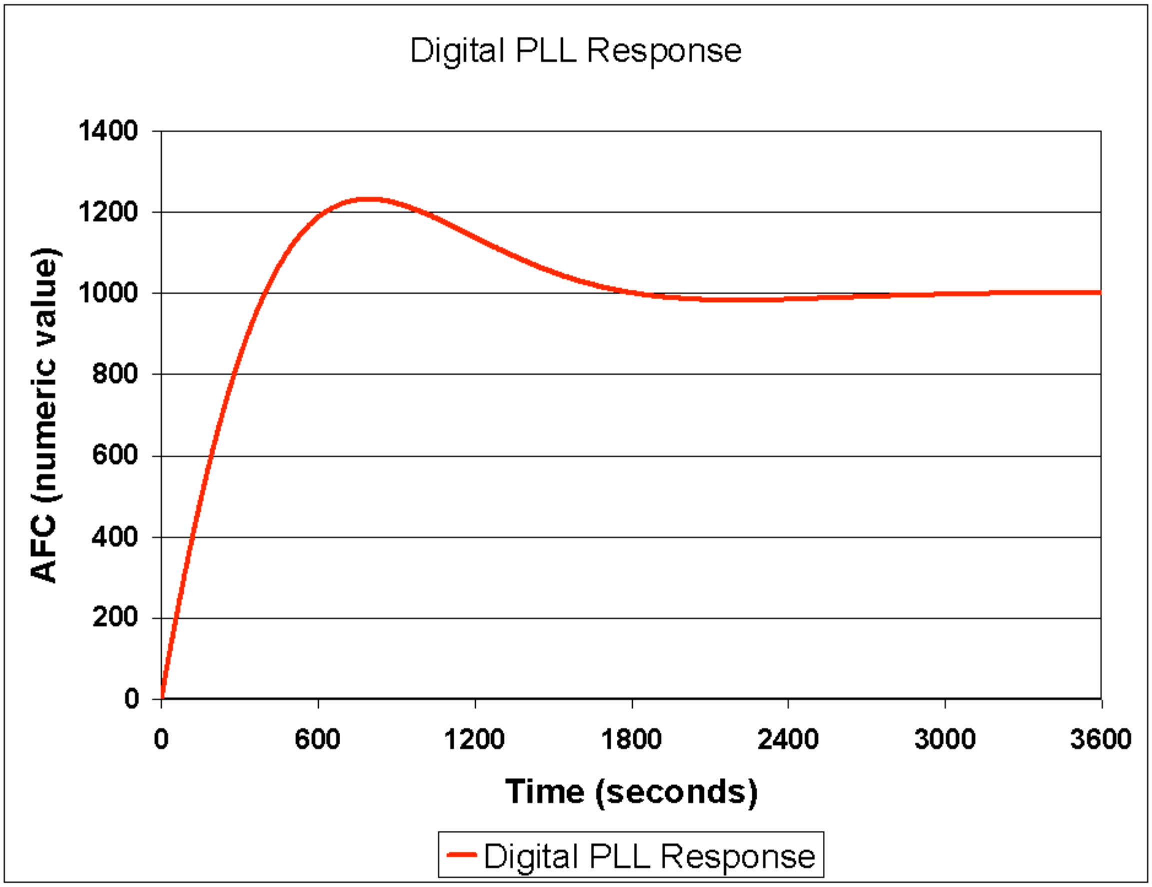

Figure 12 – Step Response of Slow DTxA Digital PLL for Clock Frequency Smoothing

The architecture of the DTx system was set up so that transmitters from different manufacturers could be combined to form a network, and to still have all transmitters remain on frequency even if the SMPTE 310 clock is off frequency, or even if it is slewing. Different modulator designs react differently to SMPTE 310 static and dynamic clock frequency errors. So, by reducing the dynamic frequency error to a low value, and by requiring all modulators to produce a correct IF frequency even if there are SMPTE 310 clock frequency errors, all of the transmitters will stay within 1 Hz of one another.

One other design technique is applied to the distributed transmitters. Parallelism is exploited to achieve a fast frequency error measurement, shortening the intervals shown in Figure 11 above. A shorter measurement interval reduces time lag and residual frequency error in the presence of frequency slew. The design objective is to keep all transmitters within 0.5 Hz even if the input SMPTE 310 clock is slewing at its maximum of 0.028 ppm/second.

The DTxA also generates a field rate side channel, which carries data to the distributed transmitters on every field. This side channel is useful mainly for the EVSB mode. In EVSB, the VSB mode bits and/or some of the reserved bits in the field sync may change on every or any field, as the distribution of enhanced and normal packets changes dynamically. Clearly, the normal DTxP cannot normally handle this, because it normally occurs much less frequently than once per data field. Typically, the DTxP occurs once a second, which is about once every 41.3 fields.

The side channel is formed by stealing the MPEG packet error bit. This gives a gross payload capacity of 312 bits per field, since there are 312 MPEG packets per ATSC data field. Roughly half of this capacity (160 bits) is given to Reed-Solomon error correction.

At the distributed transmitters, the MPEG packet error bit is restored to its proper value (0, indicating no error).

Details

The system can be examined in detail by looking at the components of the distributed transmission packet. The DTxP has been assigned a packet identifier (PID) of 0x1FFA. This indicates an "operations and maintenance" (OMP) packet. Although future OMP packet functions may be assigned, the only one that exists at this point is the DTxP.

The OMP/DTxP is shown in Figure 13 below. It contains the usual MPEG preamble words, followed by a OMP packet type byte and 183 bytes of payload.

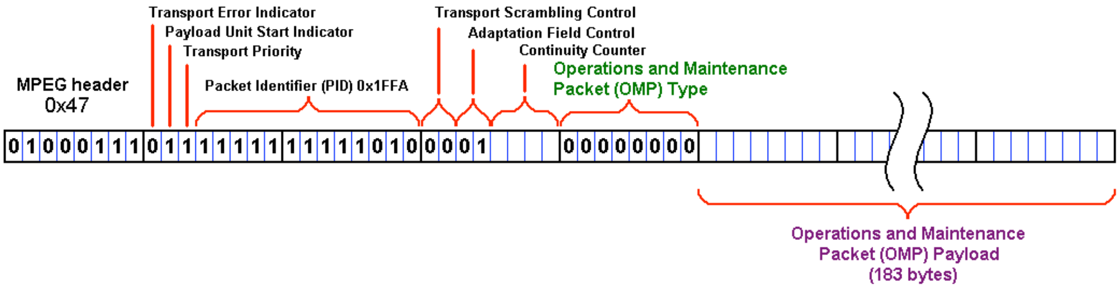

Figure 13 – Distributed Transmission Packet (DTxP)

Table 1 (taken from ATSC CS/110A) shows the payload that is carried in the distributed transmission packet or DTxP.

The first byte is reserved. The next 8 bytes contain the trellis code states. After that, there is a synchronization time stamp. The time stamp is the number of 10 MHz cycles (100 nanosecond periods) since the last 1 ppm clock tick, at the time that the first bit of the DXTP is sent out from the DTxA.

The "maximum delay" value is the approximate time delay from the output of the DTxA to the output of the distributed transmitters. Setting this value will be a function of the sum of the delays of the STL systems and the transmitters. The "maximum delay" value is adjusted for each individual distributed transmitter by another value – the "offset delay."

Table 1 — Distributed Transmission Packet (DTxP) Organization

Syntax	Bits	Format
DTx_packet () {		
Reserved	8	0xFF
For (i=0; i<12; i++) {		
Trellis_code_state	8	riuimsbfwp
}		
Synchronization_time_stamp	24	uimsbf
Maximum_delay	24	uimsbf
Network_identifier_pattern	12	0xFF FFFF FFFF FFFF FFFF
Reserved	2	‘11’
packet_number	10	uimsbf
Reserved	32	0xFFFFFFFF
Tx_group_number	8	uimsbf
for (i=0; i<16; i++) {		
tx_address	12	uimsbf
tx_identifier_level	3	uimsbf
tx_data_inhibit	1	bslbf
tx_time_offset	16	tcimsbf
tx_power	12	uipfmsbf
Reserved	4	’1111’
}		
Reserved	320	for (i=0; i<40; i++) 0xFF
DXP_ECC	160	uimsbf
}		

Where:

bslbf	bit string, left bit first.
riuimsbf	repeated, inverted, unsigned integer, most significant bit first
riuimsbfwp	repeated, inverted, unsigned integer, most significant bit first, with parity
tcimsbf	two's complement integer, msb (sign) bit first.
uimsbf	unsigned integer, most significant bit first
uipfmsbf	unsigned integer plus fraction, most significant bit first

The “network identifier pattern” is a seed value for a Kasami sequence generator that creates the RF watermark signal. Kasami sequences are created by a combination of linear feedback shift registers. They have the desirable qualities of being able to create a large number of sequences that have low cross correlation values – in other words, there is a low probability of mistaking one sequence (transmitter code) for another. The watermark code sequences are controlled by the DTxA, and the information about the codes is broadcast over the air so that test receivers can interpret it and read the network.

The “packet number” value is the pointer to the cadence sync – the MPEG sync byte associated with the ATSC frame sync. This “packet number” provides redundancy with the cadence sync. If a piece of equipment or a system will not pass the cadence sync, the “packet number” value may instead be used to locate the beginning of frame sync.

Up to this point, the data transmitted in the DTxP is used by all transmitters in the network. Next in the packet there is a group of 16 sets of values that must be addressed to individual transmitters.

First is a 12 bit transmitter address value. This allows addressing up to 4096 transmitters in a network.

Next is the RF watermark injection level. This may be set to any “bury ratio” from 21 to 39 dB below average power, in 3 dB steps, or it may be turned off completely. Bury ratios of 21 and 24 dB are intended for use during out-of-service testing periods only, as they would degrade transmitted signal to noise ratio values beyond the recommended 27 dB.

The RF watermark may be modulated with slow speed data at the ATSC field rate, or about 41 bits per second. This may be useful as a transmitter to studio (TSL) link. If this capability is used, there may be times when this modulation should be halted. The “TX data inhibit” bit tells a particular transmitter whether or not its RF watermark may be modulated with data.

The “TX time offset” value is the difference from the “maximum delay” of emission time for a particular transmitter. This value, different for every transmitter, is how the network timing is adjusted.

“TX power” is a value that can be used to adjust slave transmitter power. Implementation of this feature is optional.

Hardware and Field Experience

WPSX in State College, Pennsylvania is the first station to implement distributed transmission. The main problem is terrain shielding. The WPSX analog signal is on channel 3, and their digital allocation is channel 15. However, the transmitter is 38 airline miles away from the city of license. Intervening mountain ridges largely block the UHF signal. WPSX considered moving the tower closer to State College but for various reasons this was not feasible. And even if it were there would still be propagation problems covering Altoona and Johnstown. WPSX concluded that distributed transmission was the best way to cover their market.

Figure 14 – A Distributed Transmission Adaptor (DTxA)

The DTxA is located at the WPSX studio. The DTxA is shown in Figure 14.

Most of the functions performed by the DTxA are done on a channel coder card,, shown in Figure 15 below.

Figure 15 – Channel Coder Card Used for Distributed Transmission

The shield boxes on the left contain line interfaces, PLLs, voltage controlled oscillators, etc. All of the digital bit manipulation and processing occurs in the ball grid array (pinless) part visible about a quarter of the way across the board from the right. Below that part are five memory chips that provide the necessary time delay (up to one second).

Figure 16 – Main WPSX (Distributed) Transmitter

The main WPSX transmitter, shown in Figure 16, operates at approximately 800 kW average ERP.

Figure 17 shows the second transmitter in the network, which produces 50 kW average ERP. This transmitter is located on a ridge just south of State College, at Pine Grove Mills, Pennsylvania.

Two more 50 kW average ERP slave transmitters are planned for Altoona and Johnstown, Pennsylvania.

Figure 17 – WPSX Distributed Transmitter at Pine Grove Mills, PA

Testing

The first tests of the DXT system were made on July 3, 2003. What follows is taken from a report made by Merrill Weiss:

"The system comprises two transmitters -- one main transmitter at Clearfield, authorized for approximately 800 kW ERP, although operating at 537 kW for the initial tests while awaiting a CP modification from the FCC; and one distributed transmitter at Pine Grove Mills, authorized for and operating at approximately 50 kW ERP. The transmitters are synchronized according to ATSC CS/110A, and the transmitter timing was adjusted for the tests as will be described shortly.

“WPSX is a public television station owned by The Pennsylvania State University, and the WPSX studio is on the Penn State campus in State College, PA. The studio location is in an area that does not receive a very good signal from the main transmitter and is intended to be served by the Pine Grove Mills distributed transmitter. Because of the terrain, the D/U ratio at the studio is very high in favor of the Pine Grove Mills transmitter. Thus, to obtain better data on system operation in areas that will be more difficult for receivers, tests were conducted at the station's STL relay site on Rattlesnake Mountain.

“The Rattlesnake Mountain site was predicted to have a D/U ratio of around 7 dB in favor of the Clearfield transmitter. Measurements at the site were close to this, with a D/U ratio of 8.3 dB obtained using a bow tie antenna at about 2_ m AGL and rotated to maximize the signal from each transmitter as it was separately operated. A similar test conducted with a reflector-Yagi antenna produced a D/U ratio of about 12 dB in favor of the Clearfield transmitter.

“The time offset between the transmitters was adjusted to put the concurrent arrival time point approximately over the Rattlesnake site. The site includes a number of communications towers, and reflections from them were evident in channel impulse response measurements. The strongest reflections were very close to the main impulse and were on the order of 10 dB down. Reflections up to 39 us, the limit of the instrument used, were noted on the Pine Grove Mills signal, although at levels on the order of 40 dB down.

“With both transmitters operating and the bow tie antenna in use, a 4th generation receiver produced video output signals as the antenna was rotated through 360 degrees. (Audio was not checked for these tests, but there is no reason to expect that it would not have been properly recovered.) With the reflector-Yagi antenna aimed to favor the Pine Grove Mills transmitter, thus reducing the D/U ratio to about 1.6 dB, there was immediate lock on all four virtual channels in the stream. There were some rare video blinks, but this seemed to be due to low signal levels from a splitter that was producing more loss than expected. There was not time to obtain another splitter or to install an amplifier.

“With both transmitters operating and the bow tie antenna in use, a pre-1st generation receiver only locked on the signal intermittently, more or less as expected. When the reflector-Yagi was aimed to favor the Clearfield transmitter, thereby increasing the D/U ratio, the pre-1st generation receiver operated correctly, again as expected.“

Distributed Transmission and E-VSB

From a transmitter’s point of view, E-VSB is a messy business. Two different input streams have to be aligned, interleaved, and transmitted. The presence of two physical inputs means special hardware. Moreover, data needs to be located at specific, controlled places within a data frame. This means special firmware.

The ATSC distributed transmission ad hoc group recognized that E-VSB and distributed transmission may need to be applied simultaneously. The ad hoc group also recognized that there is a possible synergy between E-VSB and distributed transmission. There are some significant similarities between E-VSB and distributed transmission. Specifically, certain data needs to be transmitted in certain places within the data frames in both distributed transmission and in E-VSB.

The result is that an E-VSB transmitter is identical to a distributed transmission slave. The same firmware and software that allows a distributed transmission slave to insert frame syncs and data packets at controlled locations makes it possible to use the same design to transmit E-VSB.

The only thing that needs to change is the DTxA. For E-VSB, we might call it a EVSBA (enhanced vestigial sideband adaptor). The EVSBA would take in the two streams and form distributed transmission packets that would tell the slave transmitters where to put the frame syncs with respect to the data packets. And as a bonus, it would also tell the slave transmitters what the trellis coder states are and what the system timing should be. That way, the system could easily be made to transmit E-VSB from distributed transmitters.

So, the bottom line is that if you have a DTx slave transmitter, you already have an E-VSB transmitter. The only thing that needs to change is the DTxA.

Distributed Translators

In some areas, particularly in the west, there are many translators on the air. However, even in relatively unpopulated areas, there simply are not enough channels available to provide DTV translator service at the same level where NTSC translator service already exists.

In a traditional translator system, many channels may be used in addition to the originating channel. This is not an efficient use of spectrum. During the transition period, there are probably not enough channels in mostly western “translator territories” to accommodate both NTSC and ATSC using current spectrally wasteful techniques.

A more efficient use of spectrum would be to use just ***one*** additional channel for the translators, by applying distributed transmission technology.

The distributed transmission system is based on the use of STLs to transmit the modified SMPTE 310 data to ***all*** of the transmitters in a network. This is acceptable and even desirable for many applications. But translator systems are intended to receive signals off the air, and to retransmit them on another channel. Requiring “translators” to have STL systems would seem redundant – why require the same signal to be delivered to a translator site *twice*?

Operators of translator systems would like to be able to take a signal off the air, and use that signal to drive a network of synchronous distributed translators.

As originally conceived, the distributed transmission system did not support off-air operation. However, some new ideas have been developed which can make distributed translators possible. We will be discussing these new concepts in this paper. These methods will require some slight modifications to the ATSC candidate standard.

At this point we should review the process for generating a DTxP and processing it at a distributed transmitter.

First, the distributed transmission packet is generated by the DTxA. The DTxA contains a model of the channel coding process. The DTxP contains all of the system information, including the trellis coder states derived from the coding model, the system timing information, watermark codes and injection levels, etc. The only problem here concerns the trellis codes. It is not possible to know what the trellis codes are until the data including the DTxP goes through the trellis coder model. So, it is not possible to include the trellis coder states in the DTxP before trellis coding it.

To deal with this causality problem, the DTxP is only partially formed, with the trellis coder states omitted. Dummy data is inserted into the fields reserved for the trellis coder states. Then the packet in this form is passed through the channel coding model.

At that point, the trellis coder states become known, and then they are inserted into the DTxP.

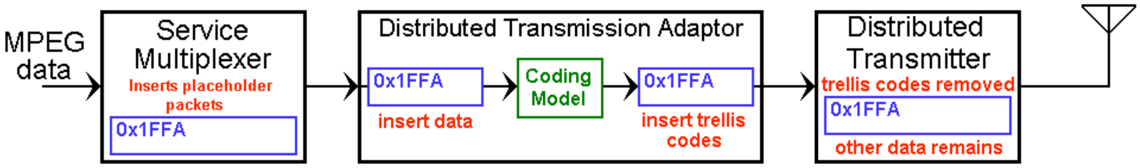

Figure 18 – Life Cycle of a Distributed Transmission Packet

If a distributed transmitter were to simply extract these codes and then pass the unmodified DTxP through its channel coder, then its channel coder would no longer match the model in the DTxA. The reason is that the slave transmitter and the DTxA's coding model are not operating on the same data. The DTxA is operating on a packet with dummy trellis codes, and the slave transmitter is operating on a packet with actual trellis codes.

Every bit must be identical in both the DTxA and the slave transmitters. If just one bit is off, the slave transmitter will turn into a "jammer" because it will not be generating the same symbol sequence.

So, each slave transmitter must remove the trellis coder data from the DTxP, and replace it with the dummy data that was the basis of the channel coder model used in the DTxA. When the DTxP is transmitted over the air by the distributed transmitters, it will not contain any trellis code information. The "life cycle" of a DTxP is shown in Figure 18.

So, the bottom line is that when the DTxP is put on the air, the trellis coder state information is lost. (However, the DTxP that goes on the air contains all of the other information about the network, including the watermark sequences, the transmitter timings, etc. This information is useful to test and measurement equipment.)

That means that downstream transmitters do not have access to the trellis codes. Without the trellis codes, such downstream transmitters would all turn into mutual jammers, since they would be transmitting different symbol streams from one another.

Fortunately, there are solutions to this problem. The nature of the solution depends on the complexity of the translator system.

If the translator system is simple (meaning single hop translators only), then the solution is simple. If the translator system is complicated (multiple hop) then the solution is a little bit more complicated.

Let’s consider the simple translator system first. In the simple system, we have a “main” transmitter and a number of translators in surrounding areas. All of the translator sites can receive the signal from the main transmitter.

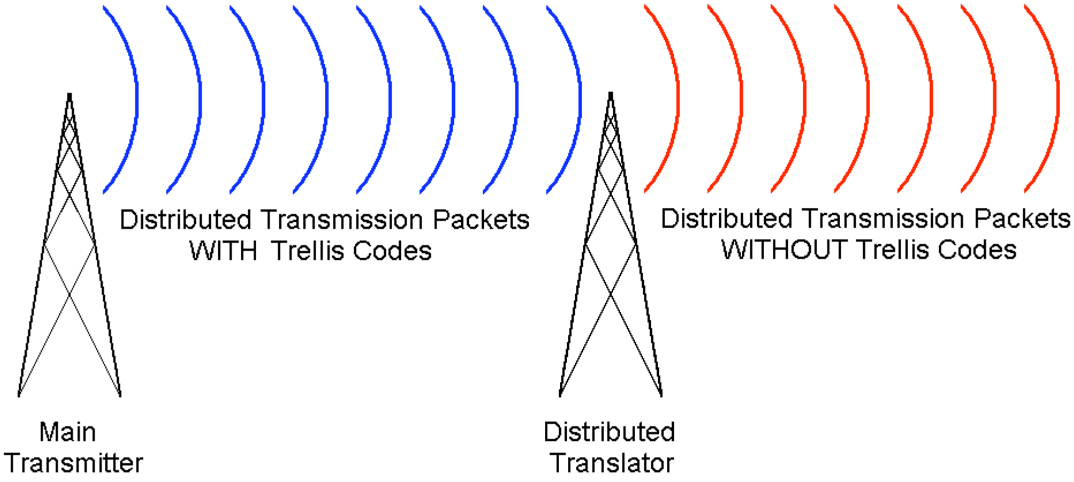

Figure 19 – Single Tier Distributed Translator Network

In this example, shown in Figure 19, assume that the originating station is on channel 8. All of the translators might use channel 41, for example. The channel 8 signal would serve as the “STL” to all of the channel 41 transmitters. There is only one main transmitter, and it alone is on its channel, so it does not need to synchronize with anything. The DTxP, carrying the trellis codes etc. would be transmitted over the air by the channel 8 transmitter. Since the channel 8 transmitter is operating as an ordinary ATSC transmitter, it will not change the data in the DTxP. The channel 41 distributed transmitters would receive the DTxPs, decode them, and adjust their timings and trellis coders etc. in response to the commands embedded in the DTxP. The channel 8 and channel 41 transmitters will operate with different trellis codes, different symbol sequences, and different timings. But, since they are on different channels, this is not a problem. The DTxP will control only the channel 41 transmitters, and the channel 8 main transmitter will operate autonomously.

One other thing happens when we use a normal ATSC transmitter as an STL: the inversion of the SMPTE 310 MPEG sync bytes, which indicates cadence sync, is lost. This happens because MPEG sync is replaced with segment sync. However, cadence sync, normally indicated by an inversion of the MPEG sync byte, is transmitted redundantly in the distributed transmission system. A pointer to the frame sync insertion point is included within the DTxP (this is the "packet number" value in Table 1) . So, all distributed translators will instead refer their frame sync insertion points using that pointer. Although the MPEG sync inversion is available once per ATSC data frame, the frequency of the pointer will be the same as the DTxP insertion rate. This is typically once per second, but it may be set to any rate desired by the system operator.

When the DTxP is received from the channel 8 transmitter, it contains all of the trellis coder state data. But when the DTxP is transmitted by the channel 41 transmitters, that data must be removed so that the trellis coders in each distributed transmitter can follow the states of the encoding model present in the DXTA.

So, the trellis coder states are no longer available at the outputs of the channel 41 transmitters. This simple technique is therefore limited to "one hop."

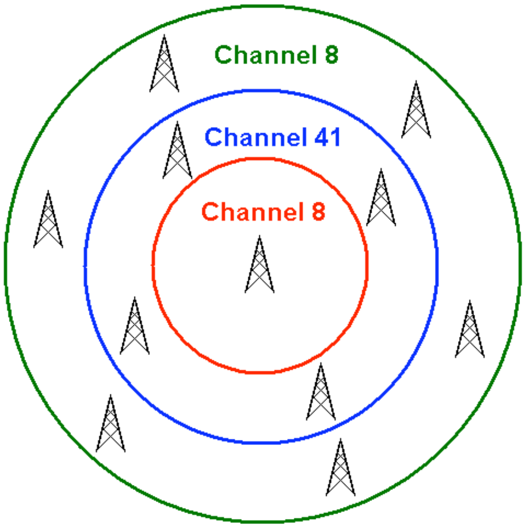

Figure 20 – Distributed Translator Network Using Only Two Channels

If we want to have a multiple hop distributed translator network, how can we get trellis codes (along with the other necessary data) sent, over the air, to downstream distributed transmitters? This concept is shown in Figure 20. There are several ways to do this, and we will discuss one method in detail.

One way to do this would be to use a special receiver (which does not yet exist) which extracts trellis codes and frame sync insertion points off the air. Since trellis codes are the "first line of defense" against distortions, they are the first to become corrupted. So, off-air trellis codes would have to be validated by having the special receiver test for trellis decoder errors. This approach would have some fragility, and would require a special receiver.

Another way of extending a distributed translator network to multiple rings of transmitters would be to insert several layers of DTxPs. In effect, a number of DTxAs could be cascaded at the studio. The first DTxA would insert DTxPs for the last (outer) ring of transmitters. The second DTxA would insert DTxPs for the next to the last tier of transmitters, etc. Each tier would have additional delay to allow for interleaver and other coding delays. And, each tier would operate with different trellis codes.

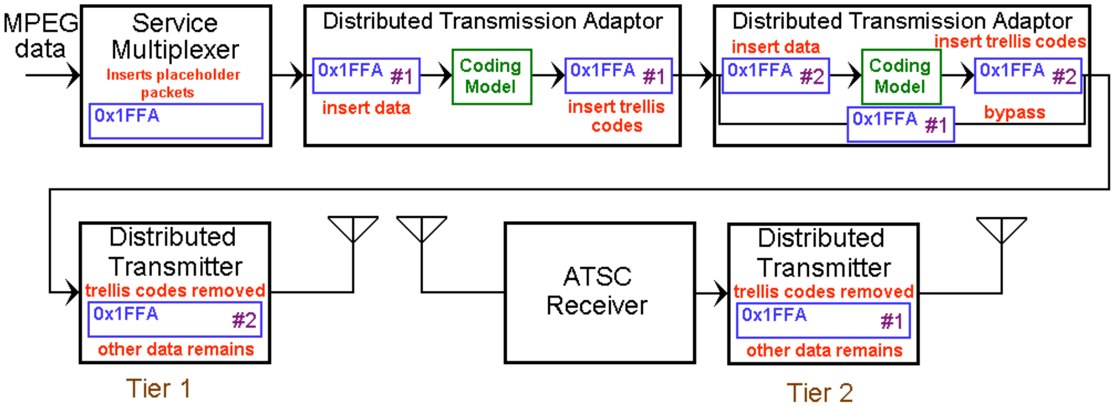

Figure 21 – Multilayer Distributed Transmission System

Figure 21 shows the layering of the DTxPs. Each DTxA inserts a layer of DTxPs. Each DTxA processes only the DTxPs corresponding to its tier of transmitters, and the other DTxPs bypass the encoding process. At each transmitter tier, the associated DTxPs are interpreted, and the embedded trellis codes are removed. One layer of DTxPs is required for each tier of transmitters.

Channel reuse is also possible. Extending our example, channel 8 could be the main transmitter, and the first tier of distributed translators would be on channel 41. A second tier of translators could receive the channel 41 signal, and retransmit on channel 8 (assuming that the signal from the second ring of translators does not interfere with the main transmitter). A third tier could be back on channel 41 again.

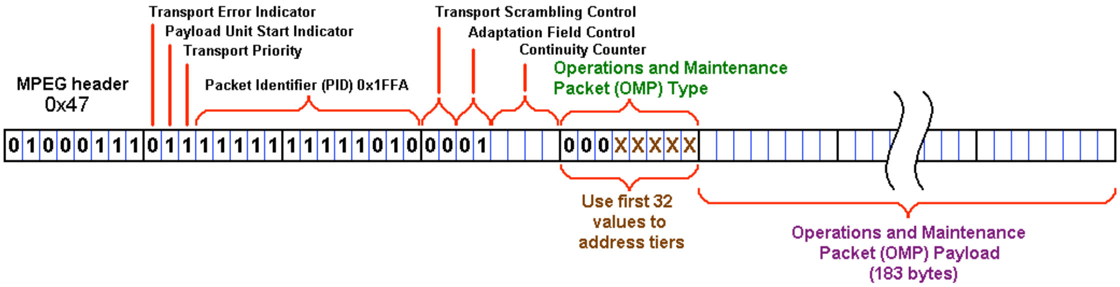

Figure 22 – Distributed Transmission Packet with the Ability to Address Multiple Tiers of Distributed Translators

Figure 22 shows how the DTxPs could be modified. The OMP Type byte could use the first 32 values (from the lower five bits) to point to a particular tier of transmitters. Alternatively, the first reserved byte shown in Table 1 could instead be used to indicate the addressed tier of distributed transmitters.

E-VSB Side Channel in Distributed Translator Systems

The distributed translator system can transmit all of the timing information, trellis codes, etc., but it cannot transmit the side channel. But, the E-VSB data, including the VSB mode bits and the field sync reserved bits, may be taken off the air and retransmitted. This will require a special receiver for the E-VSB mode that can recover and output these bits. This special receiver would only be required for E-VSB distributed translators. No special receiver is required for distributed translators with conventional ATSC signals.

There are two other methods that could be used to handle E-VSB in a distributed translator system.

First, the side channel could be changed from using the MPEG error bit to the MPEG header priority bit. The priority bit is not currently used (as far as we know) in the ATSC system. So, the side channel could possibly be transmitted over the air this way.

We know from experiments that transmitting data using the MPEG error bit definitely affects receivers with severe picture breakup. So using a station’s main transmitter as an “STL” for a translator network will not work with a side channel present in the MPEG error bit. Receivers reject MPEG packets with the error bit set.

The other way to make E-VSB work with multiple levels of distributed translators would be to move the field rate data into the DTxP’s reserved bits. Then, the DTxP transmission rate could be increased to once per field if necessary (using up 0.32% of the channel capacity per tier of distributed translators). Or, the E-VSB system could be operated such that a DTxP is only transmitted when the E-VSB packet mix ratio needs to change.

System Design

Terrain shielding makes for easy design of distributed transmission systems. And terrain shielding is often the reason for needing a distributed transmission system in the first place. With terrain shielding, there will be few areas where signal overlap is a problem.

In flat terrain areas, there may be reasons to use distributed transmission. If tower space is not available for a single full power transmitter, distributed transmission may be an alternative.

In flat terrain areas, there are some pitfalls to system design. In free space, RF travels one mile in about 5.37 microseconds. So, moving away from one transmitter and towards another on a straight line would change the equalizer timing at twice that rate, or about

10.74 microseconds per mile. If the two transmitters are both high powered, then it would not take very many miles to exceed the equalizer window in some receivers.

So this example of how not to design a distributed transmission system suggests some basic system design rules. The objective of these design rules is to avoid creating large areas with both (1) low D/U ratios, and (2) high timing displacements. Terrain shielding helps in this effort. Other tools include directional transmitting antennas, beam tilt, choice of antenna height and power level, and locating interference zones in unpopulated areas.

The worst possible design of a distributed transmission system would be to put up at least two high powered transmitters, tens of miles apart, with tall towers above flat terrain, spanning a populated area.

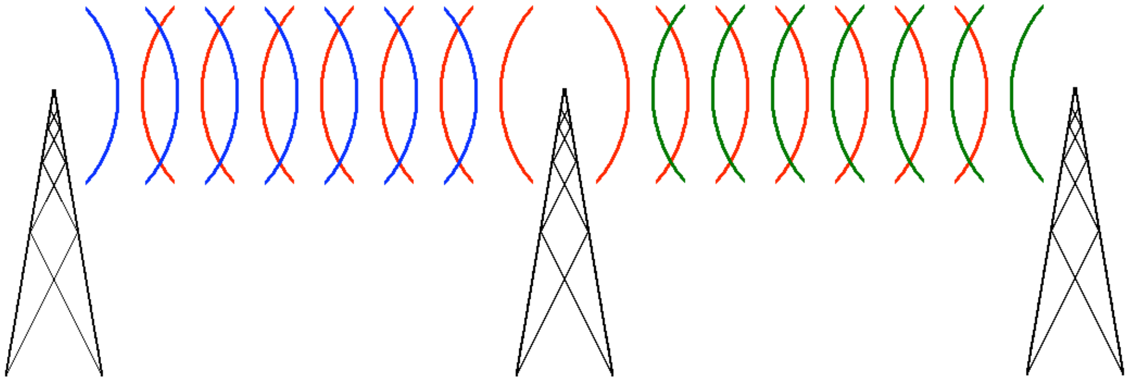

Figure 23 – A Poor Flat Earth Distributed Transmission System Design

This would ensure that there would be large areas with low D/U ratios with over 100 microseconds of delay spread. This kind of system is shown in Figure 23. But, even with such a poor design, almost everybody could obtain reception by using a directional antenna.

A better design of a "flat earth" system is shown in Figure 24. In this example, slave transmitters feed directional antennas pointing away from a "main" transmitter, and the timing of the slaves is adjusted to be coincident on a collinear line drawn away from the two transmitters.

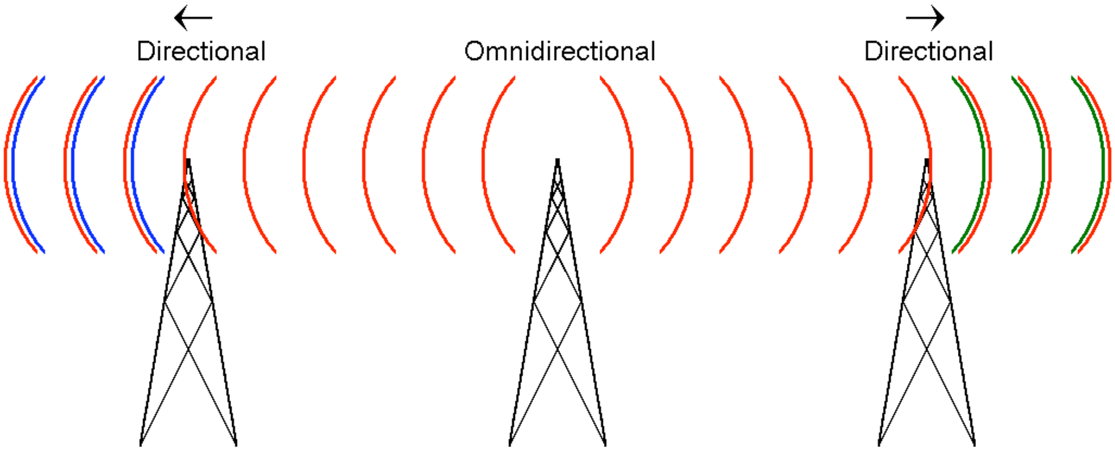

Figure 24 – A Better Flat Earth Distributed Transmission System Design

Conclusions

The ATSC system has turned out to be quite flexible, and has lent itself well to the creation of single frequency networks, also known as distributed transmission.

It is also possible to apply this technology in distributed translator systems.

Such systems can be designed to work even with legacy receivers. Newer receivers with improved adaptive equalizers will work better in a distributed transmission environment, and such receivers may ease DTx system design.

Distributed transmission is another option in our digital television implementation toolbox.

Bibliography

[1] Weiss, S. Merrill, "Distributed Transmission Systems — Overcoming the Limitations of DTV Transmission," NAB 2003 proceedings, pp 263-279, 2003.

[2] Hershberger , David L.; Pugh, Michael W., "Implementation of the ATSC Distributed Transmission System", NAB 2003 proceedings, pp 280-289, 2003.

[3] Mattsson, Anders; Weiss, S. Merrill; Simon, Mike; Hershberger, David; Wu, Yiyan; Wang, Xianbin, "Transmitter Identification Techniques for Distributed Transmission Networks," 53rd Annual IEEE Broadcast Symposium, 2003.

[4] Advanced Television Systems Committee, Candidate Standard CS/110A, 2003.